

Security for VAX Systems

digital

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by Digital Equipment Corporation or its affiliated companies.

The following are trademarks of Digital Equipment Corporation:

DEC	DECCL	Scholar Plus
ALL-IN-1	MASSEUS	SES
DEC	MicroVAX	ULTRIX
DECmate	Packnet	UNILIS
DECmail	PDP	VAX
DECnet-VAX	P/OS	VAXcluster
DECserver	Professional	VAX DBMS
DECsystem-10	Rainbow	VAX RDC
DECsystem-2	RIS/VMS	VAX 8600
DECUS	RSTS	VMS
DECwriter	RSX	VT
DESNC	Scholar	Work Processor

MS-DOS is a registered trademark of Microsoft Corporation. UNIX is a registered trademark of American Telephone and Telegraph Company.

Copyright © 1989 Digital Equipment Corporation. All Rights Reserved.
Printed in U.S.A.

**Security for
VAX Systems**

EC-G0027-31

Table of Contents

Chapter 1 • Overview

Why Computer Security Is Important to You	1-1
Security in Today's Computing Environment	1-1
Customizing Your Computer Security	1-1
Types of Computer Security Problems	1-2
The Reference Monitor Concept	1-3
Security Features of VMS	1-4
Guide to VMS System Security	1-4
Optional File Security for VMS	1-5
Network Security	1-5
Recovery Services and Disaster Planning	1-5

Chapter 2 • System Security

Security Capabilities of VMS	2-1
Access Control Features	2-2
User Identification Codes	2-3
Logging In	2-3
Passwords	2-5
System Operation	2-5
Auditing Features	2-7
DECnet Security	2-7
Enhanced Ethernet Security System	2-9
Security Models	2-10
VAX Software Products	2-10
ALLAN-1	2-10
Rdb Database Security	2-11
VAX DBMS Database Security	2-11

Chapter 3 • Security for the System Manager

Introduction	3-1
Elements of Computer Security	3-2
Threats to Computer Security	3-2
Establishing User Accounts	3-2
Password Management	3-4
Login Options	3-6

iv • Table of Contents

The Secure Terminal Path	3-6
Controlling the Number of Retries on Dial-Up Logins	3-8
Controlling Break-In Detection and Evasion	3-8
Authorizing Usage	3-9
Educating the New User	3-10
Guarding Against Trojan Horses and Computer Viruses	3-10
Disk Maintenance	3-11
Disk Scavenging Countermeasures	3-12
Summary of Disk Scavenging Prevention Techniques	3-13
Auditing with Security Alarms	3-13
Other Audit Data	3-15
When Your System Is Under Attack	3-15
Security Auditing	3-17
Handling a Security Breach	3-17

Chapter 4 • File Protection

Introduction	4-1
Standard UIC-Based Protection	4-2
UICs and Protection	4-2
Specifying UICs	4-2
How UIC-Based Protection Controls Access	4-2
Protection Code Syntax	4-3
How Privileges Affect Protection	4-4
How the System Interprets a Protection Code	4-5
How the System Determines Access Types for Objects	4-5
Access to Disk Files	4-5
Access to Directory Files	4-6
Access to Volumes	4-7
Access to Global Sections	4-7
Access to Devices	4-7
Access to Logical Name Tables	4-7
Access to Queues	4-8
Establishing and Changing Volume Protection	4-8
Establishing and Changing Directory Protection	4-9
Establishing and Changing UIC-Based Protection	4-9
Access Control Lists (ACLs)	4-10
Managing Access Control Lists	4-13
Disk Scavenging Countermeasures	4-14
Managing Your Files for Optimum Security	4-15

Chapter 5 • Network Security

Security for a DECnet Node	3-1
The Reference Monitor in a Network	3-1
Establishing Subject Correspondence	3-2
Proxy Accounts	3-2
Setting Up Proxy Logins	3-2
Specifying Authorizations	3-3
Protecting Communications	3-3
Enhanced Ethernet Security System	3-4
The Security Modem	3-9
Security for Packet-Switched Data Networks	3-10
Terminal Servers	3-11
DECnet-VAX Accounts	3-15
The DECnet-VAX Database	3-16
Network Usage	3-16
Security Concerns on a Cluster	3-16
Using DECnet Between Cluster Nodes	3-17
Summary	3-17

Chapter 6 • Auditing

Introduction	6-1
Auditing with Security Alarms	6-1
Audio Reduction Facility	6-4

Chapter 7 • Security for the User

Introduction	7-1
Logging In	7-1
Passwords	7-2
If You Have Multiple Accounts on Different Systems	7-4
Summary of Password Guidelines	7-5
Causes of Login Failures	7-5
System Password Failures	7-6
Login Class Restrictions	7-6
Shift Restrictions	7-6
Dial-Up Login Failures	7-6
Break-In Evasion	7-7
Symptoms and Causes of Login Failures	7-7
Network Security Considerations for Users	7-8
Logging Off the System	7-9

Appendix A • Running VMS in a C2 Environment

The Trusted Computing Base	A-1
Protecting the TCB	A-1
Individual Accountability	A-2
Object Protection and Reuse	A-2
Protection of the Audit Trail	A-3
Auditing Actions of a System Operator or Administrator	A-3
Documentation	A-3
Physical Security	A-4
Configuration Guidelines	A-4

Appendix B • Managing Classified Data

Glossary

Figures

2-1 VMS Reference Monitor	2-6
2-2 Network Reference Monitor	2-8
3-1 Secure Terminal Path	3-7
5-1 User Account Routing	5-3
5-2 Proxy Account Routing	5-4
5-3 Security Controls for Ethernet	5-7
5-4 ISO Reference Model for Open Systems Interconnection (OSI)	5-8
6-1 Security Alerts System	6-2
7-1 Password Encryption	7-4

Chapter 1 • Overview

• Why Computer Security Is Important to You

The data processed on your computer system is among your organization's most vital assets. Such data may include confidential financial results, trade secrets regarding products, classified government information, or fund transfers. Compromise of this information by unauthorized disclosure, fraud, sabotage, or natural disaster may have serious financial and operational consequences. Thus, protecting the integrity of data and programs should be a top priority for any organization that relies on computers in pursuing its goals.

• Security in Today's Computing Environment

In the 1950's and 1960's, the computer resources of most businesses or government agencies consisted of large mainframe computers. Housed in well-protected, limited-access areas, these resources were relatively easy to defend from most forms of compromise.

Today, the proliferation of small personal computers, the phenomenon of distributed computing, and the increasing reliance on computer networks present unique security challenges. The Digital Equipment Corporation has responded to these challenges by building a high degree of security into its VMS operating system, and by offering a range of services, hardware, and software that address specific security needs.

• Customizing Your Computer Security

Ultimately, the security of your data and communications, whether computer-based or not, depends on the effectiveness of your overall security program. Security features accomplish little unless consistently, conscientiously, and intelligently applied. Many security breaches involve legitimate users exercising authorized access and privileges, rather than malevolent outside forces breaking into your system.

In considering the material in this handbook, bear in mind that not all security procedures described may be desirable for your organization or site. The Digital Equipment Corporation's approach to security permits great flexibility, so features may be implemented in ways that precisely fulfill your particular needs.

A related point to consider is that trade-offs are inherent in protecting information. Typically, an extremely secure system is less user-friendly than a less secure system, and more time may be required to perform certain operations. In striking a balance between security and efficiency, it is necessary to realistically assess your requirements. For example, the classified data processed by a government installation generally demands a higher order of protection than the administrative data of a university.

Another consideration is that not all your organization's data demands the same degree of protection. Among your challenges is to provide appropriate security procedures for the kinds of information processed on your system.

• Types of Computer Security Problems

Breaches of security generally fall into one of three categories, depending on the technical skills and resources of the user. These categories are:

- User irresponsibility
- User probing
- User penetration

User irresponsibility involves misuse of authorized access, and demands little or no technical skill on the part of the user. For example, a user may abuse his authorized access to certain files by making a copy of a sensitive file and selling it to a business competitor. Such occurrences are amenable to technological solutions only to a limited extent, primarily through auditing of authorized actions. The knowledge that an auditing capability exists may work as a deterrent, and auditing can help identify an offender after a breach is detected.

Note that problems in this category are more likely to develop when authorized use far exceeds the legitimate needs of the user.

User probing refers to deliberate attempts by insiders or outsiders to exploit weaknesses in system controls for the purpose of gaining unauthorized access to the system.

When access is achieved, the prober may "steal" information by reading it, or he may cause direct damage by modifying or deleting material. The activities of computer hackers may represent probing. The prober's motivation is often the intellectual challenge of beating the system. Because probing represents an invasion of confidential information, it should be taken seriously even in the apparent absence of malicious intent. VMS provides a number of features to combat probing.

User *penetration* implies an attempt to break through security controls for the purpose of gaining control of the system.

Penetration always suggests malicious intent, and often involves sophisticated methods. Military installations and other sensitive government sites are customary targets for penetration attempts. The security features of VMS reduce, but do not entirely eliminate, the possibility of successful penetration.

• The Reference Monitor Concept

Digital's approach to system security is based on the reference monitor concept.

The *reference monitor concept* views computer system elements as subjects, objects, authorization database, audit trail, and reference monitor mechanism. These elements are defined as follows:

- *Subjects* are active processes that gain access to information on behalf of people.
- *Objects* are passive repositories of information that need to be protected.
- The *authorization database* is an entity that defines the system's security requirements by revealing which subjects can access which objects.
- The *audit trail* is a record of access attempts that is specified by the authorization database.
- The *reference monitor mechanism* is an entity that enforces security rules by mediating the creation of subjects, granting subjects access to objects according to the requirements of the database, and recording events as necessary in the audit trail. Ideally, the reference monitor mechanism should:
 - Mediate every attempt by a subject to gain access to an object.
 - Provide a tamperproof database and audit trail that are thoroughly protected from unauthorized observation and modification.
 - Remain small, simple, and well-structured so that its effectiveness in enforcing security can be assured.

By monitoring the basic structure of the reference monitor model, VMS achieves a level of security that is resistant to probing and to many attempts at penetration.

• Security Features of VMS

VMS provides *discretionary access controls*, which permit individually named users to be either included or excluded from accessing a file or achieving certain forms of access (that is, READ, WRITE, EXECUTE, DELETE, CONTROL). The reference monitor also uses *access control lists (ACLs)* to protect files and other objects. VMS prevents disk scavenging (obtaining disk space that contains another user's data) by overwriting the disk space before reassigning it to another's use.

VMS can notify a security officer, a system manager, or an operator of a security event and write an *audit record* of such events. Security events may be attempts (successful or unsuccessful) at system or file access. Auditing may be specified by type of file access. VMS can also record authorized file access by privileged users. These events can be invoked selectively on a per file basis. Other events that can be audited are changes to user account files and mounts/dismounts of tapes and disks.

Login enhancements are provided in the following ways:

- Forced hangups on multiple login failures
- Break-in detection and disabling of accounts for a period of time after detection of a break-in attempt
- Automatic account expiration
- Account restrictions based on time of day and day of week
- Restrictions on type of login, for example, allowing only local logins and disabling dial-up or network access

Password vulnerability can be minimized by establishing a minimum password length and specifying a password expiration period. A system manager may also require use of a random password generator, which creates lists of nonsense words from which a user may choose a password.

• Guide to VMS System Security

The *Guide to VMS System Security*, packaged with your VMS documentation kit, includes the following information:

- Describes security features of VMS
- Offers suggestions for security management of VAXsystem clusters and networks
- Describes how to provide a National Computer Security Center (NCSC) Class C2 Trusted Computing Base (TCB) with VMS

Optional File Security for VMS

VMS has a user-callable encryption utility (VAX encryption) for file encryption that is a software implementation of the Federal Data Encryption Standard (DES). This option is under special export control.

VMS architecture also has implemented mandatory access controls on files and devices. Efforts are continuing to complete the implementation of mandatory controls required by the B level criteria of the NCSC. At present, this capability is provided by the Digital's *Security Enhancement Service*. (For a detailed description of the Security Enhancement Service, see Appendix B.)

• Network Security

Protection of information processed on a network is a complex challenge, in part because such data may be subject to interception.

Digital is pursuing several approaches to improving network security, including incorporation of encryption options into its networking products.

For local area networks, Digital has introduced an *Ethernet Security System* to assist in the implementation and management of security measures in an Ethernet LAN. This product is under special export control.

Note:

Customers outside the United States should check with their account representatives regarding availability.

Other means of enhancing network security, such as use of security modems, are discussed elsewhere in this handbook.

• Recovery Services and Disaster Planning

Threats to your computer system and data are not limited to intrusion. Major physical damage to your system and stored data may result from fire, natural disaster, sabotage, or act of war.

Optional services (available in the United States) can help restore your computer operations in the aftermath of a catastrophic event that damages or destroys your data or system.

1.5 • Overview

Digital also offers an Escrow Storage Service. It protects you from loss of third-party vendor software support by protecting software source code for your critical third-party applications and releasing the software to you if the vendor is unable or unwilling to provide maintenance. In addition, Digital has a wide range of products and services designed to protect and retain data residing on disk storage.

The following chapter discusses system security in more detail. Other chapters are devoted to the specific concerns of system managers, general users, and auditors.

Chapter 2 • System Security

• Security Capabilities of VMS

VMS is the operating system for Digital's VAX series of 32-bit minicomputers in both standalone and VAXcluster configurations. VMS affords a high degree of security by providing mechanisms for controlling access, securing the database, and creating an audit trail.

Among VMS security features are:

- A user entry and password system
- An operations log that includes provisions for alerting security staff of ongoing security breaches
- An encryption system providing a higher level of security for sensitive files

VMS runs on all VAX hardware and is compatible with the MicroVAX, thereby facilitating both upward migration and networking. This operating system has been evaluated by the National Computer Security Center (NCSC) and in August, 1986 was rated at the C2 level. VMS thus became the first minicomputer general operating system to receive a security rating of C2.

One of the great advantages of VMS is the flexibility with which its security features can be implemented. For example, file protection features can be invoked to protect all, some, or none of your files. Similarly, provisions can be made for use of general passwords, locked passwords, or no passwords. The use log may be ignored, used only for record keeping, or employed as a security tool.

By selectively using the various security features of VMS, it is possible to provide a high degree of security for sensitive data while eliminating or reducing control procedures for less critical information. After analyzing your specific security needs, you can use the security capabilities of VMS to control access, discourage random abuse, and deter serious compromise.

In short, VMS provides as much control as can reasonably be expected of a general purpose operating system with only minor inconvenience for legitimate users.

• Access Control Features

The essence of system security is controlling access to data.

An authorized user, by definition, is entitled to access at least some data on the system. Of course, this access may be severely limited. For example, a user may be authorized to access only one file on the system, and he may be authorized only to read that file. Another user may be authorized to change the contents of the file, as well as to read it. A third user may be authorized to read the file, change its contents, and then delete it.

These three individuals are all authorized users, but the nature of their authorized use is sharply defined, and varies from one to another. A fourth user may not be authorized to access the system at all, and would more properly be referred to as an intruder.

The primary control mechanisms that VMS uses to restrict access are *access control lists (ACLs)* and *user identification codes (UICs)*. The ACL uses *identifiers* to specify users. The three types of identifiers, are:

- *UIC identifiers*, which refer to each user on the system
- *General identifiers*, which are specified by the security manager to refer to groups of users
- *System-defined identifiers*, which describe users based on their use of the system

System-defined identifiers are automatically defined by the operating system when access rights are created during system installation. The categories are BATCH, NETWORK, INTERACTIVE, LOCAL, DIALUP, and REMOTE.

A user is automatically assigned one of these identifiers during login, and the VMS login software adds the appropriate identifier to the process rights list. In addition, an ACL may be constructed to restrict access to a dataset to particular individuals or categories of users.

An ACL consists of one or more *access control entries (ACEs)*. The three security-related ACEs are:

- An *identifier ACE*, which controls the type of access permitted to an individual or group (Access types are READ, WRITE, EXECUTE, DELETE, CONTROL, and NONE.)
- A *default protection ACE*, which defines the default protection for directory files only
- A *security alarm ACE*, which provides an alarm message to alert managers to possible security breaches when designated forms of access (READ, WRITE, EXECUTE, DELETE, or CONTROL) are accomplished or attempted

User Identification Codes

Each system user has a UIC and each system object also has an associated UIC, which is defined as the UIC of its owner. In addition, a protection code relates a user to permitted types of access. Access to objects is thus controlled by the relationship between the UIC of the user and the UIC of the object.

All users seeking access to the system fall into one or more of the following categories:

- **System**, including all users who have system privileges
- **Owner**, meaning the user who created the object to be accessed and therefore has the same UIC
- **Group**, who are all users with the same group number in their UICs as the object's owner
- **World**, which includes all users

The protection code uses these categories to permit or deny READ, WRITE, EXECUTE, or DELETE access.

When a user logs in, the identifiers that are in his rights database are copied into a rights list that becomes part of that person's process. (The process includes the user's UIC and the system-defined identifiers.) VMS uses the rights list to perform all protection checks.

Although login is a simple procedure from an authorized user's viewpoint, it is far from a simple matter of granting or denying access from the viewpoint of the system. Login sets in motion a sophisticated and highly complex process that determines the kind of access (if any) permitted to the user, what objects he can access, and what he can and cannot do with those objects.

Logging In

Users gain access to the system by logging in. Login is an initial screening process that permits the system to clear a user attempting access. The login process usually requires both a username and password so the system can check authorization and impose restrictions. Different classes of logins accommodate all possible modes of access. Logins fall into two categories: interactive and noninteractive.

An *interactive login* is accomplished when the user follows system prompts that appear on his terminal screen. For example, a local login procedure might involve specifying the system on which you have an account when prompted to "SELECT SYSTEM", and then typing your name at the "Username:" prompt and your password at the "Password:" prompt. In this case, you are providing information requested by the system, and the system is responding to your input.

To access the system you must correctly input all three items. In other words, you must specify a system that can be accessed on the terminal you are using, the username must be authorized to access that system, and the password must match that of the username. An apt analogy is that of opening a combination lock by dialing the required digits in the proper order.

A *noninteractive login* is performed by the system without user/system interaction.

Types of logins are LOCAL, DIAL-UP, REMOTE, NETWORK, BATCH, and SUBPROCESS.

- A *local login* is executed by a user who is directly connected to the system.
- A *dial-up login* requires the same procedure as a local login. In this case, the terminal/system connection is accomplished by means of a modem.
- A *remote login* involves the same procedure as a local login. This difference is that a remote login is performed to a node over a network.

Note:

Local, dial-up, and remote logins are all interactive.

- A *network login* involves user access of a file stored in a directory on another node, or performing a network task on a remote node (assuming that both nodes are on the same network).
- A *batch login* is a noninteractive login procedure accomplished when a batch process initiated by a user starts to run.
- A *subprocess login* is another noninteractive login that results from a user executing either a specific process form of a command or a system service.

Login Types	Interactive	Noninteractive
LOCAL	X	
DIAL-UP	X	
REMOTE	X	
NETWORK		X
BATCH		X
SUBPROCESS		X

Passwords

Passwords are used as part of the login process. In many environments, a user password suffices to access the system.

As an added security precaution, a system password may be required before the user password. In maximum security environments, two user passwords, entered successively, may be required for access.

Passwords may be either selected by the user or automatically generated. Passwords may consist of letters, numbers, or a combination of both. Security is enhanced by requiring a minimum character length for user-generated passwords, and by creating passwords resembling English in the automatic generation mode.

In addition, VMS can limit the number of password entry attempts to discourage attempts to guess a password.

All passwords are stored in a one-way encrypted form. One-way encryption ensures that no individual, including the system manager, can determine an actual user password. If a user forgets his password, the system manager can assign a temporary password to allow the user to log in to his account. After login, the user should select a different password. The system manager can force the user to enter a new password by "expiring" the temporary password.

System Operation

VMS security features are based on the *reference monitor concept*, using the reference monitor as the central security point for subjects, objects, reference monitor database, and security auditing.

When a user logs in, VMS creates a process with the identity of the user. The process gains access to information as the agent for the user in the system. Process creation and access to information by processes are the critical mechanisms by which the operating system enforces security.

Because process creation raises a number of potential security problems, many security features of VMS are focused on this area. When attempting to log in, a user must provide both a username (which is given to the resulting process) and a password. The file of encrypted user passwords becomes part of the reference monitor's database. VMS provides additional password security by storing encrypted passwords in a file that is normally excluded from general access.

After a process is created, VMS assigns it a user identification code (UIC), which corresponds to the UIC of the user creating the process. The UIC also indicates the group to which the user belongs. Other information may be attached to the process, for example, the affiliation of the process's owner with various other groups.

The most basic objects in the reference monitor concept are files and directories. These are protected in a number of ways from unauthorized access, while a variety of mechanisms allow for controlled sharing of data.

Objects other than files and directories can be used to store sensitive data. Among these objects are sections, mailboxes, logical names, and event flag clusters.

In the reference monitor concept, each subject's authorization to gain access to an object is contained in the authorization database. The database is distributed and stored in association with objects that must be protected. For example, the authorization data for a directory is included in the header for that directory.

VMS permits great flexibility in the implementation of its security features.

Most objects are protected on a UIC basis, which specifies whether access is allowed or denied to processes acting on behalf of system management, the object owner, other members of the same UIC group as the owner, and all other users.




Figure 2-1 VMS Reference Monitor

In addition to UIC-based protection, files and directories can be placed under control of *access control lists (ACLs)*, which list users or groups allowed or denied access to particular files and directories. ACLs specify sharing on the basis of UIC and other identifiers associated with a process. For example, it is possible to specify that a file cannot be read by a process connected to a dial-up line, thereby preventing an otherwise authorized user at a remote location from accessing that file over the telephone line.

• Auditing Features

Detection of possible security breaches is an important element in system security.

VMS allows a terminal to be designated as an audit alarm console on which all auditable events may be displayed. Such events may include login failures and successful or unsuccessful attempts to access sensitive files. Events may be audited at the discretion of users and system managers. For example, the owner of a particular file can create an ACL entry that requires an audit of any attempt to access that file.

• DECnet Security

As might be expected, the basic reference monitor concept is similar whether a single system or a network must be protected.

Implementation of the reference monitor concept over a network is more complex. Accessing an object across a network involves two systems: one containing the subject seeking access and the other containing the object.

Because the reference monitors of each computer can deal only with subjects and objects on their individual systems, an access attempt across a network involves a phantom object and a real subject on one system, and a phantom subject and a real object on the other.

Of the various means VMS provides to establish correspondence between a subject or process on a source node and another on a target node, the use of proxy accounts offers distinct advantages. This option requires the target reference monitor to maintain a table of source subjects (specified by user name and node name) and the corresponding target (or local) user names.

Each request from a subject on a source node is then mapped into the creation of a subject representing the corresponding target user. This mechanism offers the explicit control associated with username/password control while adequately protecting passwords.




Figure 2.2 Network Reference Monitor

The security of network operations depends on the ability of source and target reference monitor mechanisms to communicate in a private and authentic manner. For obvious reasons, it is critical that an intruder be unable to observe passwords or mimic a source node that has been granted proxy access. The most effective methods of achieving such protection are physical protection, such as that afforded by coaxials, and encryption.

• Enhanced Ethernet Security System

The *Enhanced Ethernet Security System*, composed of *DESNC security network controllers* and *VAX/VMS Ethernet security manager software*, provides security management and control in an Ethernet LAN. DESNC controllers for Ethernet local area networks upgrade security through:

- Authentication of Ethernet nodes
- Enforcement of a mandatory access policy among Ethernet nodes
- Data protection through encryption and integrity controls

DESNC controllers provide transparent cryptographic security at a level consistent with the National Institute of Standards and Technology (NIST) Data Encryption Standard (DES).

With the Enhanced Ethernet Security System, you can logically separate different classes of nodes in a LAN. For example, development and production systems can be physically on the same Ethernet bus without the ability to send Ethernet packets to each other. PCs and workstations can be similarly managed with the Enhanced Ethernet Security System.

The DESNC controller is a store-and-forward device providing real-time cryptographic processing of Ethernet frame (messages) over an Ethernet LAN (local area network).

Decryption restores the data to its original form through a client node that takes advantage of encryption/decryption services.

A DESNC controller can accommodate any combination of workstations, servers, or VAX processors. Client nodes supported by the DESNC controller can run operating system software that includes VMS, RIX, and ULTRIX, as well as industry standard operating systems such as UNIX* and MS-DOS.†

* UNIX is a registered trademark of the American Telephone and Telegraph Company.

† MS-DOS is a registered trademark of the Microsoft Corporation.

VAX/KDC is a layered VMS product that serves as the central authority for managing the DESNC controllers and enforcing security policy for an Ethernet LAN or extended LAN. Multiple VAX/KDC nodes may be used to improve the availability of the networked DESNC controllers.

• Security Modems

Security breaches may be attempted by intruders using modems. Security modems help combat this problem. Digital's *Scholar Plus Modem* features an access security system that protects data with a system of callback security and password verification.

Unauthorized access is prevented by means of modem parameters, dial memory, and callback memory locations. Mandatory password access prevents unauthorized modification when the modem is unattended. The callback security feature screens incoming calls before allowing access to the host system, and provides an audit trail of access attempts.

The Scholar Plus modem enforces security with little inconvenience to a remote user, who simply inputs the correct password and/or telephone number for validation by the modem. Upon validation, the Scholar Plus performs a callback or password, and the user is on-line and ready to transmit data.

The Scholar Plus automatically adjusts to the operating speed of Digital's terminals for all speeds between 1200 and 9600 bits per second.

• VAX Software Products

ALL-IN-1

ALL-IN-1 software provides a working environment in which users have ready access to personal and common resources. For example, all ALL-IN-1 users have access to the time management menu, which is a common resource. But all appointments the user creates through that menu are stored in a file in his directory, which is a personal resource accessible to him alone.

ALL-IN-1 runs on the VMS operating system, which provides a hierarchical file storage framework and various protection capabilities for directories, files, and other system objects. ALL-IN-1 uses the directory framework and file protection features to:

- Retain a basic structural integrity while allowing extensive customization
- Provide a personal environment and a sense of ownership for each user while maintaining an efficiently managed, shared environment
- Prevent accidents that can result from resources being accessed by unauthorized users

In customizing the ALL-IN-1 directory and resource protection framework, VMS commands can be used to set UICs, specify user privileges, and set protection on a resource.

Additional protection mechanisms are contained in the ALL-IN-1 User Profile, which controls user access to resources.

Rdb/DBMS Security

Rdb/VMS relational database software provides a security mechanism to protect the database from unauthorized browsing or modification. This mechanism applies specifically to Rdb/VMS operations. It builds upon existing VMS security features to provide a more finely grained set of access controls to mediate READ and WRITE access.

Rdb/VMS security depends on access control lists (ACLs) attached to databases and relations. ACLs determine which users can access database entities and what operations they can perform. Specifically, Rdb/VMS security can limit access to data manipulation statements, data definition statements, and utility statements.

The three steps in defining protection for a database or relation are:

- Deciding access rights to be granted to specific users and creating a set of *access control entries (ACEs)*
- Arranging ACEs in proper order
- Constructing the ACL

In most respects the security functions of Rdb/VMS software work in a manner that is essentially similar to those of VMS. As in VMS, the identifiers are *UIC identifiers*, *general identifiers*, and *system-defined identifiers*.

VAX DBMS Database Security

The VAX DBMS system uses three primary features to enhance database security:

- A *security scheme*, which specifies a set of data manipulation language (DML) access permissions
- A *User Execution List (UEL)*, which maps a user to a security scheme
- *Command Language Identification Lists (CALs)*, which identify users authorized to access VAX DBMS operator utilities

2.12 • System Security

Use of security schemas is optional for each database. They are used to restrict DML access to the database for some users. For example, confidential data, such as that relating to employee salaries, may be made available only to specific users.

Because a database typically contains information to which all users must have at least retrieval access, managers should protect database entities in an appropriate and selective manner. Such decisions are usually based on an identification of the various types of users and an assessment of the kinds of access each needs.

Chapter 3 • Security for the System Manager

• Introduction

The system manager's primary challenge is to implement established security policy in ways consistent with the needs of his organization or site. This involves realistically assessing these needs, and then addressing them with appropriate and effective procedures.

While it is tempting to aspire to a state of maximum security, a comprehensive effort to protect all data may be neither necessary nor desirable. Digital's approach lends itself to customizing, and the wise manager takes full advantage of this capability.

In achieving the desired level of security, the system manager inevitably spends considerable time interacting with users.

Training users is of prime importance. It imparts knowledge of specific security procedures and instills an awareness of security issues. Most users cooperate in security measures if the manager explains how they work and why they are important. In the absence of such explanations, some users may adopt a dangerously casual attitude toward security.

Many security breaches result from a user's failure to observe elementary precautions, rather than from a sophisticated effort to defeat system defenses. For example, a user who leaves his terminal unattended while logged in has left the door open to the system, rendering it vulnerable to a potential intruder. Fostering a positive attitude toward security is one of the most important functions of the security manager.

To put the problem in perspective, it is useful to note that human errors and omissions account for approximately 50-60% of annual computer related losses. (Deliberate transgressions by authorized users represent about 10%, and intrusions by outsiders only 1-3%.) Clearly, the security manager who focuses on compromise from without may fail to properly guard against dangers from within.

• Elements of Computer Security

An effective security program includes three elements.

- First, a security policy must be formulated that specifies who is allowed to do what and under what circumstances they are allowed to do it.
- Next, a physical and organizational environment must be created that affects where and how the security policy will be implemented.
- Finally, computer system mechanisms and controls must be used to enforce the security policy.

• Threats to Computer Security

The three threats to computer security discussed in Chapter 1 of this handbook are:

- User irresponsibility resulting in breaches caused by accident or deliberate fraud
- Probing for ways to gain access to information on the system
- Penetration by a malicious intruder seeking to gain control of the system

The best defenses against user irresponsibility are application controls on user activity, administrative controls, and audit trails.

Probers may be deterred by system controls. Successful probing often results from lapses such as use of a too short or too obvious password or a failure to adequately protect files.

Penetration is the most difficult form of attack to thwart. The most successful methods of defending against penetration are a security kernel-based operating system that averts malicious attempts to bypass system controls, and encryption of communication lines. These methods generally exceed both the needs and resources of most organizations, and are addressed in this handbook only to a limited extent.

• Establishing User Accounts

Information security doesn't mean locking up your data and hiding the key. To be useful, information must be accessed and shared as well as protected. The security features of VMS facilitate the kind of fine control that keeps information flowing while defending it from unauthorized access.

One of the key steps in this process is establishment of user accounts. As a system manager, it is your responsibility to decide who gets to do what with objects on the system.

One approach is assigning users to user identification code (UIC) groups. Groups are usually formed according to function, because users performing similar or identical functions may be assumed to require access to the same information. Members of the accounting department might constitute one group, and members of the marketing department might constitute another. In this example, payroll information, which is confidential, would be available to the accounting department but not to the marketing department.

However, UIC-based protection has only limited usefulness, because in reality members of different groups usually must share at least some of the same data. This sharing could be accomplished by making the information available to the WORLD category of user, but this would strip the data of all protection because any user on the system could access the data.

A better approach to controlled access is the establishment of access control lists (ACLs). ACLs can be assigned to files, and define the access allowed to groups and individuals. For example, the ACL for the PAYROLL.DAT file might read as follows:

```
(IDENTIFIER=DON,ACCESS=READ-WRITE-EXECUTE-DELETE)
(IDENTIFIER=JONES,ACCESS=READ-WRITE-EXECUTE-DELETE)
(IDENTIFIER=SMITH,ACCESS=READ-WRITE-EXECUTE-DELETE)
(IDENTIFIER=JOHNSON,ACCESS=READ)
(IDENTIFIER=THOMAS,ACCESS=READ)
```

As a system manager, you can use the AUTHORIZE Command to simplify this ACL by creating the general identifier PAYROLL for Don, Jones, and Smith, resulting in the following:

```
(IDENTIFIER=PAYROLL,ACCESS=READ-WRITE-EXECUTE-DELETE)
(IDENTIFIER=JOHNSON,ACCESS=READ)
(IDENTIFIER=THOMAS,ACCESS=READ)
```

This is a good approach because VMS can process shorter ACLs more rapidly, and because it simplifies matters when personnel changes occur. For example, if Don is transferred to another department and is replaced by Johnson, you simply remove Don's user account file (UAF) record, depriving him of the identifier. You would then grant Johnson the right to hold it.

ACLs permit great flexibility in controlling access to data, and are best used when objects demand a highly refined degree of protection.

It is also possible to create identifiers based on type of system access, that is, LOCAL, DIAL-UP, REMOTE, INTERACTIVE, NETWORK, and BATCH. For example, the access control entry (IDENTIFIER=THOMAS+LOCAL,ACCESS=READ) specifies that Thomas can

read the file, but only if he is logged on a terminal directly connected to the system. He would not, for example, be able to access the file from his home terminal using a modem.

After deciding the names of the identifiers you want on your system and the users you want to hold them, use the VMS AUTHORIZ utility to make these associations known to the system.

• Password Management

Password management is your first chance to screen out the undesirables, that is, unauthorized users. Most commercial sites use passwords. A maximum security site often uses primary and secondary passwords, and perhaps a system password.

The first requirement for a password is to be unguessable. This eliminates many obvious choices, such as the name of your dog, your nickname, or the date of your wedding anniversary.

When you open an account for a user, you give the user a username (usually his surname) and an initial password (which the system requires him to change when he logs in for the first time).

As stated previously, passwords are stored by the system in an encrypted form. Because the encryption is one-way, no one can deduce a password from its encoded equivalent. Even a system manager cannot retrieve a user's password. If the user forgets his password, the system manager can set a temporary one to allow the user to enter his account and choose another password.

Additional system passwords may be used to restrict access to certain terminals. A *system password* is a single system-wide password applied to selected terminals.

A user who logs on a terminal or dial-up port protected with a system password must enter the password before the system identifies itself and asks for "Username" and "Password". Examples of such terminals include those that:

- Use dial-ups or public data networks for access
- Are publicly accessible (and not tightly secured), such as computer labs at universities
- Located in remote areas and not frequently inspected
- Intended only as spare devices
- Reserved for security or privileged operations

System passwords are primarily useful in high security environments to thwart unauthorized access while concealing from an intruder the name of the system he is attempting to probe.

Primary and secondary passwords are used when security considerations demand the presence of two users to access the system. Because this procedure involves visual contact with the users, it greatly minimizes the risk of unauthorized access. Obviously, this makes sense at a high security site, but the inconvenience of the procedure limits its practical application in less rigorous environments.

You may also consider increasing the default minimum password length and/or expiration period.

The minimum length requirement discourages adoption of easily guessed passwords, and mandatory expiration forces the user to change passwords often enough to frustrate a potential intruder. The default life of a primary or secondary password is six months. In general, users with access to critical files should have shorter password expiration periods.

System passwords have an unlimited lifetime, but should be changed at intervals of six months or less. If a user with access to the system password leaves, the system password should be changed immediately.

Password security can be enhanced by requiring use of the random password generator. This mechanism generates password choices that resemble English words, but are passwords not found in a dictionary. Use of the random password generator makes sense if you are protecting privileged accounts, or accounts with access to especially sensitive data.

It might also be wise to implement this measure if you have reason to believe your users may be careless in choosing passwords or maintaining their privacy.

Use of the random password generator can be imposed on any or all of your users. For example, an analysis of your site's particular security requirements may prompt you to demand that all privileged users use the random password generator.

For added security, you might further specify that generated passwords have a minimum length of eight characters. (Use of the random password generator may be established as a precondition for granting of privileges.)

A special situation arises when your system needs servicing by a Digital Field Service representative.

It is your responsibility to issue a unique password to the Field Service representative each time service is required. You should also disable the FIELD account when the service has been completed.

The Field Service representative is responsible for maintaining the confidentiality of the password, and he will neither change the password nor perform any tasks beyond those associated with the servicing of your system.

Default passwords are changed and the accounts disabled during the installation of your system. Thereafter, it is your responsibility to manage the FIELD account and issue a password to the Field Service representative.

• Login Options

You should consider the amount of information displayed at login. As a system manager, you can control the appearance of announcement, "welcome", last login, and new mail messages.

Note:

It may be inadvisable to "welcome" a user to the system, because an intruder may use this phrasing as a legal defense to justify unauthorized access.

In a tight security environment, you are well-advised to be conservative with messages that appear before a user has successfully logged in. Specifically, you may wish to avoid giving clues that may reveal the identity of the operating system. At a maximum security site, prudence may dictate dispensing with an announcement message.

Similarly, consider the advisability of altering the welcome message, which by default states the operating system and version number, and often states the node as well. Control is achieved by writing your own message, which may warn against unauthorized access, or printing a blank line in lieu of the information you wish to conceal. It is also possible to disable the welcome announcement for individual users.

Last login messages may likewise be suppressed. In this case, you may enable it for specific users.

If you really want to be stingy with your information, you can suppress the new mail announcement, which reveals the number of new electronic mail messages directed to an account. This is especially appropriate if you've chosen to deny a user access to the VMS MAIL utility.

• The Secure Terminal Path

The *secure terminal path* is a VMS capability that is designed to thwart "password grabbers". These are programs that steal passwords by mimicking an empty video screen, a screen showing that the system has just been initialized after a crash, or a screen showing a logout.

The "password grabber" seeks to lure the user into logging in, at which time the program ascertains his password (to the benefit of the intruding programmer).

When the password is revealed, the screen often indicates a login failure. This misleads the user into assuming he has mistyped his password, and conceals the deception that has just been perpetrated.

To combat this problem, you can invoke the secure terminal path, which terminates any currently executing process before the start of a login. The secure terminal path is invoked on a terminal by terminal basis with the following DCL command:

```
SET TERMINAL/PERMANENT/SECURE/DISCONNECT
```

To initiate a login on a terminal protected in this way, the user must press the BREAK key and then the RETURN key. The login is then accomplished in the usual manner.

Note that the secure terminal path is useful only when the terminal is connected directly to the system. It is not an option when the user is logging in over the network.




Figure 3-1 Secure Terminal Path

• Controlling the Number of Retries on Dial-Up Logins

Limiting the number of retry attempts for dial-up logins hinders an intruder from gaining access by entering a known username and then using a small computer programmed to enter every word in the dictionary as a password.

Limiting retries does not present a problem for the user who makes a few typos. He is granted a grace period and a number of chances to enter his password before being disconnected. The default allows for three retries with a twenty second interval between each attempt. You may change the parameters as you wish. However, the values may not be applied selectively, but will affect every user on the system who has dial-up access.

• Controlling Break-In Detection and Evasion

Limiting retries on dial-ups is sound policy, but it's not enough to thwart a determined intruder, who won't be inordinately discouraged by the prospect of repeated redialling. VMS provides a number of additional mechanisms to address this problem.

VMS automatically defines a threshold count that indicates a possible break-in attempt when five unsuccessful login attempts have been made to an account within a specified interval. Login failures in this case are those caused by invalid password entries. Failures must come from one of three specific sources:

1. A specific terminal and a specific valid username (This is the default.)
2. A specific remote node and a specific remote username
3. The username of the creator of a detached process

By default, five attempted logins from one of these sources are permitted. However, you can adjust that value as you wish.

Evasive action consists of refusing additional login attempts for a specified period of time. That means, if a terminal is the source, no one can log on that terminal while using the username that is now under suspicion. If the source is a node, the specific remote user is unable to log in.

This mechanism discourages attempts to learn a correct password through trial and error. However, this protection is not absolute. An astute intruder may spread his break-in attempts over a prolonged period, thus evading the detection mechanism. He may also change terminals, nodes, and/or usernames frequently enough to avoid suspicion. It is obvious that the detection feature complicates his job.

In lieu of disabling access to an account for a specified period after a suspected break-in, you may choose to disable the account until you manually intervene.

• Authorizing Usage

As you authorize users, consider what restrictions make sense in balancing the needs of their jobs with the demands of security. You may restrict a user to certain devices, commands, privileges, working times, and/or modes of operation, for example, BATCH, DIALUP, REMOTE, NETWORK, LOCAL, or INTERACTIVE. You can also impose an expiration period.

In some situations, you may want to create captive accounts for certain users. Consider this option when:

- Permitting unskilled or semiskilled workers to perform routine computer tasks
- Running batch operations during unsupervised periods
- Running applications programs with information that you want to keep private

A *captive account* (also known as a *dummy account*) allows limited access to the system, usually through a specialized login procedure. In addition to limiting the activities of the user, a captive account customarily denies access to the DCL command level. When creating a captive account, it may be advisable to disable the welcome announcement and electronic mail. (The VMS MAIL utility has the potential to spawn other processes.)

In establishing a captive account, two password options are available if you do not want the user to control the password:

- No password, which is specified by the AUTHORIZE qualifier /NOPASSWORD
- A locked password (allows only the security manager to change it), which is specified by the /FLAG=LOCKPWD qualifier

Locked passwords are generally preferable to open captive accounts (those with no password). If you assign a locked password, give that password to all users of the captive account. Your application may require you to impose additional AUTHORIZE qualifiers on the account, such as /NODIALUP, to restrict modes of operation. You can also restrict periods of the day and days of the week when the process can be run.

You may want to use captive accounts to restrict particular users to certain applications. In this case, each user should have a separate account with separate passwords under the control of the user.

• Educating the New User

Most users cooperate in observing good security practices if you explain how these practices work and why they're important. Some cases of user probing result from users "fooling around" or experimenting with the system while oblivious that they're doing anything wrong. To head this off explain clearly, and right from the beginning, precisely what the user is and is not authorized to do.

Another advantage of carefully educating the user is to anticipate and disarm an ignorance defense in the aftermath of problem behaviors.

Consider the use of a standardized form in educating users. This ensures that all users start off with the same basic information. Do not, however, simply hand the form to the user. Rather, sit down with him and go over it point by point. Augment your presentation when possible with demonstrations of salient features.

• Guarding Against Trojan Horses and Computer Viruses

A *Trojan horse* is a hostile software program that performs a useful function but also has a hidden, destructive purpose. The destructive functions the Trojan horse may perform include deleting files, creating new user accounts, and displaying false information. A *computer virus* is a special form of Trojan horse that replicates itself from user to user and system to system.

A Trojan horse or virus may be unwittingly introduced into the system by a user who received it from a "friend" or colleague. After it has gained access to the operating system, the virus can copy itself onto any program on any disk until the computer is turned off and restarted with an "uninfected" disk.

Viruses can remain dormant for a specified period of time, or their hostile action may be triggered by a specific event, such as removal of the programmer's name from a company's payroll. The virus may be programmed solely to wreak havoc (for example, by deleting data) or to seek information residing in files and other system objects.

Although there is no absolute defense against viruses, you can take these steps to minimize the threat they pose:

- Back up important data and program disks.
- Boot up your computer from a hard disk, or from one original, write-protected operating system disk.
- Turn your computer off and boot up again from your "safe" disk before switching to a different program disk.
- Never leave your computer on and unattended if others have access to it.
- Be extremely wary in using any software lent or given to you by others. It may be sound policy never to use such software.
- Ensure compliance with software change control procedures to protect against an insider's introduction of a Trojan horse.
- Warn your privileged users (for example, system managers and system programmers) to avoid executing (or even deleting) unknown software while they possess privileges.
- Regularly monitor the modification dates of critical software such as system files. Review the User Authorization File to ensure that all entries have been authorized. Check the file protection settings for unauthorized changes or weaknesses (for example, READ or WRITE access to the World category).
- Monitor abnormal system activity or degraded performance that may indicate a change in the security state.
- Audit and monitor the use of special privileges or system services.

• Disk Maintenance

Disk maintenance measures fall into four categories:

1. Physical security for disks
2. Backups of disks
3. Physical security for the backups
4. Retrieval of files from backups

In addition to good physical security, you can take two steps to protect data residing on disks. First, establish in advance effective backup procedures to facilitate recoveries when files are deleted. Secondly, adhere to a firm policy of never giving backup media to a user.

This second point is crucial. The user can make his own copy from the backup media and then analyze the contents at his leisure, thereby rendering futile all your efforts at physical security.

• Disk Scavenging Countermeasures

Disk scavenging, which is reading "deleted" data from a disk, may be a problem at medium security to high security sites. In this context, we are primarily concerned with an unauthorized person gleaning information from a file that has been deleted or purged.

It is important to remember that deleting a file simply removes its header from the directory, and allows it to be eventually overwritten by new material. Meanwhile, it can be scavenged.

Your first line of defense is physical security, but physical security is not enough. Using UIC-based volume protection to restrict access to disks containing sensitive information is a big step in the right direction. VMS also provides two other mechanisms to combat this problem: data erasing and highwater marking.

Inclusion of the ERASE qualifier on the DELETE and PURGE commands overwrites the entire file location with an erasure pattern of zeros. You may encourage users to specify their DCL commands as DELETE/ERASE or PURGE/ERASE. To more efficiently practice and enforce *erase-on-delete*, enable the feature for the entire volume to ensure that all files are erased when deleted.

INITIALIZE/ERASE is an important DCL command. It erases all data on a volume, rendering it safe to remove from the facility or recycle. It also simultaneously enables *erase-on-delete* for the volume at volume initialization time.

VMS provides an optional *data security erase* pattern of zeros that is applied during a single WRITE operation over a given area. You can also specify a random pattern of zeros and/or multiple overwrites instead of the default. Multiple overwrites may be helpful in thwarting technology that detects and reads faint residual magnetic impressions. If there is danger of a disk being stolen and scanned by forces possessing such advanced equipment, multiple overwriting is sound policy.

Although the details are beyond the scope of this handbook, information is available about customizing the data security erasure pattern to fit your needs.

Another weapon against disk scavenging is *highwater marking*, which tracks the furthest extent to which a file has been written and prohibits user access beyond that point.

The VMS implementation of the principle known as *erase-on-allocate* achieves a similar end result by a slightly different means. When a file is to be created or extended, VMS determines how much disk space is required and applies the erase pattern to the extent of that space. The file is then written into the area over the erase pattern. If a user attempts to read beyond that extent, he finds only the erase pattern. In that way, the file is protected to its highwater mark.

Summary of Disk Scavenging Prevention Techniques

Disk scavenging can be discouraged in the following ways:

- Good physical security
- UIC-based volume protection
- Use of the ERASE qualifier when files are deleted or purged
- Use of default highwater marking on disks containing sensitive data

Some of these measures involve trade-offs which you should evaluate with reference to your particular security requirements. For example, erasures are time consuming. Multiple erasures are even more time consuming. You might consider, therefore, applying erase security only to key disks or specific users.

Remember that disk scavenging is still possible if disks are removed from the premises and worked over by experts armed with advanced technology. Physical security in high-risk environments assumes paramount importance in minimizing this threat. It may be advisable to demand that all disks remain on the premises, and that old disks be shredded before being discarded.

• Auditing with Security Alarms

Security alarms are messages sent to the security operator's terminal when certain events occur. Alarms can help you monitor break-in attempts and other undesirable activity at your site. For example, you can enable an alarm that tips you off when a user's UAF record changes.

Use of security alarms requires that you:

- Choose events to be audited
- Enable the audit features for those events
- Enable a security operator terminal
- Use the alarm information

3.14 • Security for the System Manager

The events that can be audited are:

- Selected types of access to files and global sections
- Events requested by an ACL on a file or global section
- Use of privilege to access files and global sections
- Installation of images
- Logins, logouts, and break-in attempts
- Modifications to the system and network IAF
- Changes to system and user passwords
- Modifications to the rights database
- Changes to system and user passwords
- Execution of the SET AUDIT command
- Volume mounts and dismounts

You select events to be audited by specifying one or more keywords to the ENABLE qualifier of the SET AUDIT Command. Selection of events to be audited requires restraint. If you choose too many different events, you can expect a lot of alarm messages. As a result, you will find it impossible to follow up on all of them in an effective manner. Furthermore, you will eventually cease to be particularly alarmed by the alarms.

In enabling a terminal as an alarm console, choose one that is located in a secure location and provides hard copy output. You may have one such terminal, or several. An alarm console is not required if you really don't need it. In that case, security alarms still go into the operator log file.

The information included in an alarm message depends on the type of event. All messages, however, have the same four elements:

1. The OPCOM heading, which includes the date and time the alarm was sent
2. The type of alarm event
3. The date and time the alarm event occurred
4. The perpetrator of the event, as identified by the username and process identification (PID)

• Other Audit Data

In addition to security alarms, VMS provides additional data that is useful in tracking system activity.

The system accounting log contains records of all system job terminations, including all interactive, batch, and network jobs, as well as print jobs and other process terminations. Optionally, activations of all or selected images may also be included in the accounting log.

Most network operations, including mail delivery and access to files from remote network nodes, initiate a network server job for which a log file is created. You may be able to use the NETSERVER.LOG in tracking events initiated over a network.

• When Your System Is Under Attack

An important element of system security is the ability to recognize when your system is under attack.

Because implementation of prevention features can be burdensome, it may in some cases be appropriate to actually wait for indications of trouble before taking precautions. The obvious risks in this approach may make sense only in a low security to medium security environment.

Even if you have put into place all or many of the security features discussed in this handbook, you must still monitor your system carefully for signs of attack. You must also be prepared to deal with them if they occur.

Initial indications of trouble may come from user observations, personal observations, and ongoing auditing applications.

User observations may include:

- Missing files
- Unexplained forms of last login messages
- A sudden inability to log in
- An inexplicable occurrence of break-in evasion measures
- A report from the SHOW USERS command indicating he is logged into another terminal
- Appearance of a disconnected job message for a process he never initiated
- Existence in his directory of software which he did not write
- Unexplained changes in protection or ownership of files
- Unrequested listings generated under his username
- Sudden reduction in the availability of resources such as dial-up lines

Any of the preceding observations could be indications of serious trouble. It is your job to investigate such complaints to determine their validity.

If the report turns out to be unfounded, it may be that the user is confused. Be careful to avoid insulting the user, because tactlessness on your part may foster unwillingness to report other problems in the future. Use the opportunity to impart a better understanding of how the system works. If the user seems irretrievably befuddled by the intricacies of the system, you should be careful in assigning him certain types of access or privileges.

Although users are a valuable source of information about possible security breaches, it is your responsibility to keep your eyes open for signs of trouble. Indications you should look for include:

- Appearance on the SHOW USERS report of a user that you know could not be logged in at the moment
- An unexplained change in system load
- Missing media or program listings
- Evidence of tampering with your locked file cabinet
- Strange software in your system-executable image library
- Images in the SHOW SYSTEM report
- Unexplainable authorization of usernames
- Unexplainable authorization of PROXY users
- Unusual expenditure of processing time
- Unexplained batch jobs on batch queues
- Unexpected device allocations
- Personnel problems, for example, high turnover and bad morale
- Normal processing activity at unusual hours
- Changes in UIC-based protection or ACLs on critical files

Any one of the above indications warrants investigation.

VMS provides a number of surveillance features to help you keep an eye on your system.

You can use the Accounting utility to check for problems before they get out of hand. Watch for the following:

- Unfamiliar usernames
- Unfamiliar patterns of use

- Use of an abnormal quantity of resources
- Unfamiliar sources of logins, such as network nodes or terminals

• Security Auditing

A variety of alarms can be enabled with SET AUDIT.

- For example, auditing for login failures (LOGFAIL) and multiple failures, which are identified as break-in attempts (BREAK-IN), helps detect probing.
- Auditing for LOGIN is an effective method of tracking system use and specifically indicating which accounts are being accessed.
- Enabling FILE=FAILURE auditing, which records failed file access attempts, is effective in detecting probes.
- Protection of critical files is enhanced by application of ACL-based file access auditing to detect WRITE access. In some cases, it may be prudent to audit all types of access.
- In addition, consider auditing file access involving use of privilege.

• Handling a Security Breach

A break-in usually sets in motion a four-phase sequence of events. Those phases are detection, identification, prevention, and repair. The details of the scenario depend on whether or not the breach is successful.

Let's assume you've discovered that someone has been trying to guess passwords or is browsing files. You might have been tipped off to this situation by a user report of unexplained login failures, your own observation of unusual system activity, or by audit alarms triggered by file protection violations.

If you have enabled file auditing, identifying the file browser is fairly straightforward. Things get more complicated if the file browsing has been initiated from another node on a network. In that case, you must inspect the *File Access Listener (FAL)* logs corresponding to the time of the violations, and coordinate your investigation with the system manager at the remote node.

Identifying the password guesser is more difficult, especially if he's working from a remote source using a dial-up line. Handling the situation may force a trade-off of identification versus prevention. If you opt for identification, you may have to allow continued break-in activity while you track down the malefactor.

Three ways to fight password guessing are:

1. Enforce good password security. This may involve compulsory use of the random password generator.
2. Enable system passwords on the points of entry. If you already have a system password, change it.
3. Enable auditing of successful logins to attempt to catch the event if the intruder gets in.

Three ways to fight file browsing are:

1. If you identify the file browser, discuss the event with him. Your obvious mastery of the situation should deter future attempts.
2. Warn your users about adequate file protection. Consider random spot inspections of their file protection.
3. If file browsing across the network persists, disable the default FAL account and authorize users through proxy login accounts.

The comments above assume an unsuccessful break-in. Such an event probably has not been terribly damaging, unless the browser has stumbled across a particularly choice tidbit. In this case, the nature of the information determines the extent of the loss.

A successful break-in, on the other hand, is a traumatic event. Detection of a successful break-in may be sudden and dramatic, as in discovering that critical files have been destroyed, or it may be the gradually dawning realization that an unauthorized user has been playing with your system for an extended period of time.

Your first order of business is identifying the perpetrator. A key determination is whether he is an authorized user pulling an inside job, or an outsider. Often, this basic distinction requires careful analysis of all the facts at your disposal, including data from the audit and accounting logs, facts about the nature of the penetration, the culprit's apparent knowledge of your installation, and so on.

At this point, the trade-off mentioned earlier in this section between identification and prevention comes into play. Identification is a tough job, and may be impossible unless you allow continued break-in activity. If you decide to take this route, your best weapon is stepped-up auditing. Plant traps in procedures under your control to obtain useful information. It is also a good idea to ensure that good recovery mechanisms are in place to minimize damage if files are destroyed.

Identifying an outsider usually proves impossible, especially if he has used any form of switched communication such as a dial-up line or public data network. DECnet-VAX software may allow you to track the activity back to a source node. If the attempt involved an insider, physical surveillance may flush him out.

A switched connection break-in most often allows the perpetrator to disappear without a trace into the telephone system. Realistically, you can expect little assistance from the telephone company or from an independent long-distance telephone service. These entities operate under legal constraints that inhibit them from compromising individual privacy. You would also need help from law enforcement agencies, who might not share your sense of outrage and urgency over the violation of your system. Rather than expend energy in a fruitless quest to identify the intruder, you may be well-advised to redouble your efforts at preventing further break-ins.

The following measures may prevent further break-ins:

- Secure your authorization files.
- Change passwords.
- Clean up your system software.
- Tighten security.

Repair after a break-in involves a continuation of the four steps mentioned above, with particular emphasis on cleaning up system software.

Remember that the intruder may have introduced a Trojan horse into your system, or he may have planted trap doors to facilitate future break-in activity. Therefore, consider whether piecemeal restoration of damaged files is appropriate, or whether you must start from scratch.

Chapter 4 - File Protection

• Introduction

Much of this handbook concerns ways in which VMS controls access to the system. Many features of VMS, however, are specifically designed to control access to files on the system. VMS protects files in two ways: *UIC-based protection* and *access control lists (ACLs)*.

An interactive relationship between ACLs, the UIC-based protection code, and user privileges comes into play whenever the system evaluates a request for access. VMS determines access eligibility in the following order:

1. If the object to be accessed has an associated ACL, and if the user is entitled to access that object according to the ACL, access is granted without further screening. If the ACL neither grants nor denies access, the system determines access based on the UIC-based protection. If the ACL denies access, the system uses only the SYSTEM and OWNER fields of the UIC to determine access.
2. If the object does not have an associated ACL, the system determines access according to the UIC. Access is granted or denied on the basis of the relationship between the object's UIC and the user's UIC.
3. Under some circumstances, a user's GRPPRV, SYSPRV, READALL, and BYPASS privileges may be the basis for granting access to certain objects.

Because these points are critical to understanding how the system grants access to files, we will summarize the key facts as follows:

- ACLs are always evaluated first.
- When an ACL fails to specifically grant access, UIC-based protection is checked.
- When an ACL specifically denies access, the user may still acquire access by belonging to the SYSTEM or OWNER categories and being eligible for access through them or through possessing privileges.
- Users who possess certain privileges may be entitled to access regardless of the protection offered by the ACLs or the protection code.

• Standard UIC-Based Protection

When security managers create accounts, they take two steps affecting UIC-based protection:

1. They establish each account with a standard default protection code for all files the user creates in the initial top-level directory.
2. They designate each user as a member of a group.

The default protection code and group assignment may provide sufficient protection for your files. If more protection is needed, you can use the DCL command SET PROTECTION to provide a higher level of security.

Each user in the system has a UIC defined in the system *user authorization file* (UAF). Each object in the system also has an associated UIC, defined as the UIC of its owner, and a protection code that defines who is allowed what type of access. The relationship between the UIC of the user and the UIC of the object determines access to the object.

UICs and Protection

UIC-based protection is determined by an owner UIC and a protection code, and controls access to objects such as files, directories, and volumes. A *volume* is an entity that exists when a medium is mounted on a device. For example, a disk pack is called a volume when it is mounted on a disk drive, and a reel of magnetic tape is a volume when mounted on a magnetic tape drive.

Note that the system provides protection at the file, directory, and volume level for disk volumes, but at the volume level only for magnetic tape volumes.

Specifying UICs

A UIC may be either numeric or alphanumeric. When a DCL command requires a UIC specification, you can use either format. In numeric format, the UIC consists of a group number and a member number. An alphanumeric UIC consists of a member name and an optional a group name.

Regardless of format, the system translates the UIC into a 32-bit value representing a group number and a member number. The 32-bit numeric UIC is stored in the system rights database. The *system rights database* is a file containing information pertaining to the access rights and attributes associated with identifiers and the holders of those identifiers.

• How UIC-Based Protection Controls Access

When a user attempts to access any object, the system usually compares the user's UIC with that of the object. The only exception to this generalization is when the object is protected by an ACL that immediately grants the user access.

To understand how UIC-based protection works, you should realize that once your UIC is compared with that of the object, you fall into one or more of the following categories:

- SYSTEM**
1. All users who have the system privilege (SYSPRV)
 2. Users with low group numbers (usually system managers, security managers, system programmers, and operators)
 3. Users with the user privilege GRPPRV whose UIC group matches the group of the object's owner
 4. For files on disk volumes, users whose UIC matches the owner UIC of the volume on which the file is located
- OWNER** The user with the same UIC as the user who created the object and therefore owns it
- GROUP** All users, including the owner, who have the same group number in their UICs as the object's owner
- WORLD** All users, including those in the first three categories

Through the protection code, each of the categories can be allowed or denied READ, WRITE, EXECUTE, or DELETE access.

This list omits CONTROL access because this is never specified in the standard UIC-based protection code. However, CONTROL access can be specified in an ACL and is automatically granted to certain user categories when UIC-based protection is evaluated. CONTROL access grants the user all the privileges of the object's actual owner, and allows the user to change the protection and file characteristics just as the owner could. Thus, users in the GROUP or WORLD categories never receive CONTROL access, but those in the SYSTEM or OWNER categories always do.

The actual abilities conveyed by the different types of access vary according to the situation where they apply.

The protection code describes the type of access granted to various categories of users. For example, examine the protection code SYSTEM:RWED, OWNER:RWED, GROUP:RE, WORLD:RE. The SYSTEM and OWNER categories have READ, WRITE, EXECUTE, AND DELETE access, but the GROUP and WORLD categories have only READ and EXECUTE access.

Protection Code Syntax

When you specify a protection code, you must abbreviate the access types to one character. However, user categories can be spelled out in full or abbreviated as you prefer. Each user category is separated from its access type by a colon. When more than one user category is specified, they are separated with a comma and the entire code is enclosed in parentheses.

4.4 • File Protection

The following DCL command, which sets a protection code, illustrates syntax.

```
# SET PROTECTION=(S:RWED,OW:RWED,GRP:R,WORLD:R) DATAFILE.DAT
```

You can specify categories and access types in any order.

Remember that if you omit an access type for a category, all users in that category are denied that access. If you want to deny all access to a category, list the category without specifying access. In that case, omit the colon after the name of the category. The following example of setting a protection code denies WRITE and DELETE access to the GROUP category and denies all access to the WORLD category:

```
# SET PROTECTION=(S:RWEP,OW:RWEP,O:EP,W) DATAFILE.DAT
```

When you omit a user category from a protection code, the current access allowed that category remains unchanged. Refer for a moment to the preceding example, and then assume that protection is reset using the following DCL command:

```
# SET PROTECTION=(S:RWEL,O:RWEL) DATAFILE.DAT
```

This denies DELETE access to the SYSTEM and OWNER categories, but the GROUP category retains READ and EXECUTE access while WORLD is still denied all access.

Remember that when you set protection for magnetic tape volumes, the SYSTEM and OWNER categories always have access.

• How Privileges Affect Protection

Security managers can grant privileges to users when their accounts are created or modified. Four of these system privileges take precedence over any protection specified by UICs or ACLs. These are SYSOPR, GRPPR, READALL, and BYPASS.

SYSOPR confers the same access granted to users in the system category.

A user with GRPPR whose UIC matches the group of the owner of an object receives the same access as users in the SYSTEM category. Thus, a user with GRPPR can manage a group's files.

BYPASS confers all types of access regardless of an object's protection.

READALL confers READ and CONTROL access to the object, even if such access is denied by the ACL or UIC-based protection. In addition, the user may receive any other access granted by the protection code.

• How the System Interprets a Protection Code

To determine access to an object, the system uses the object's protection code for each user category. Categories are checked in the following sequence:

OWNER, WORLD, GROUP, SYSTEM

You can access an object as soon as the system finds a compatible category that gives you the access you have requested. If you want to deny access to a category, be sure to deny access to the outermost categories. For example, check this example to see if it really denies DELETE access to the OWNER category:

SYSTEM:RMD, OWNER:RW, GROUP:RW, WORLD:RMD.

Actually, it doesn't. Although DELETE is specifically denied to the OWNER category, that access is granted to the WORLD category, which includes all users.

• How the System Determines Access Types for Objects

As stated earlier, the specific implication of various access types depends on the specific application.

Access to Disk Files

Each file on a disk has its own protection code. The specific meanings of the various types of file access are as follows:

- **READ** access confers the right to examine (read), print, or copy the file.
- **WRITE** access lets you write to or modify the file.
- **EXECUTE** access confers the right to execute a file that contains an executable program image or DCL command procedure.
- **DELETE** access lets you delete the file.
- **CONTROL** access gives you the right to change the protection and file characteristics of the file.

Note that READ access also implies EXECUTE access, and that WRITE access permits a user to change the contents of a file, or even effectively delete it by deleting major portions. (However, WRITE access cannot be used to remove a file from the directory.) Also note that both READ and WRITE access are needed to open a file for writing because VMS does not support write-only files.

Access to Directory Files

Each directory file has a protection associated with it. This directory protection can override protection on individual files. When applied to directories, access types mean the following:

- **READ** access confers the right to read (examine) or list the directory file.
- **WRITE** access lets you write to or modify the directory file.
- **EXECUTE** access gives you the right to look up files in the directory if you specify the file name.
- **DELETE** access enables you to delete the directory file.
- **CONTROL** access lets you change the protection and file characteristics of the directory file.

Note that **READ** access implies **EXECUTE** access.

If you have **READ** access to a directory file, you can display the contents of the directory file with the DCL command **DIRECTORY**. A clever intruder may use programming to access files in a directory without using the directory in which they are listed. That's why individual files should be adequately protected.

WRITE access lets you write to the directory file. In order to create files in that directory, however, you must have both **READ** and **WRITE** access. You must also have both **READ** and **WRITE** access to rename files, or to perform any file operations that involve changes to the directory file.

EXECUTE access has a special meaning when applied to directories. **EXECUTE** access lets you use the **DIRECTORY** command to look up files that you can identify by name. You can also access files in the directory that are not protected against users in your category, unless you perform an operation that modifies the directory file. In short, **EXECUTE** access provides some, but not all, of the rights conferred by **WRITE** access.

DELETE access enables you to delete a directory file after deleting all the entries it contains. When you create a directory file with **CREATE/DIRECTORY** you do not automatically acquire **DELETE** access. To have **DELETE** access, you must use the **SET PROTECTION** command to explicitly assign **DELETE** access to the **OWNER** category.

To reiterate a point made earlier, you should protect your files at both the directory and file level.

Access to Volumes

When applied to volumes, access types have the following meanings:

- **READ** is the right to examine, print, or copy files on a volume.
- **WRITE** lets you modify or write to existing files on a volume.
- **EXECUTE** enables you to create files on the volume and write into them.
- **DELETE** confers the right to delete files on the volume.
- **CONTROL** lets you change the protection and ownership of the volume.

The **READ** access on a volume limits access to read only. Note that **EXECUTE** and **DELETE** access are not valid for magnetic tapes. Granting a category of users **WRITE** access to a tape automatically permits them to have **READ** access to the volume.

Access to Global Sections

When applied to global sections, access types have the following meanings:

- **READ** lets you map the section for read access.
- **WRITE** enables you to map the section for write access.
- **EXECUTE** confers the right to map the section for execute access. (This applies only to privileged software.)
- **CONTROL** gives you the right to change the ACL. (This applies only to PPN and page file global sections.)

Access to Devices

When applied to devices, access types have the following meanings:

- **READ** lets you issue read requests to the device.
- **WRITE** lets you issue write requests to the device.
- **CONTROL** lets you change the device ACL.

Access to Logical Name Tables

When applied to logical name tables, access types have the following meanings:

- **READ** lets you look up logical names in the table.
- **WRITE** enables you to create and delete logical names in the table.
- **DELETE** confers the right to delete the table.
- **CONTROL** allows you to change the logical name table ACL.

Access to Queues

Queue operations are restricted by UIC-based protection. Restrictions may affect the types of jobs and types of users allowed on a particular queue.

When you initialize a queue, the queue is assigned a user UIC and a protection mask. Jobs are assigned an owner UIC equal to the UIC of the process that submitted the job. Each operation that is performed on a queue or a job in a queue is checked against the owner UIC, the protection of the queue and the job, and the privileges of the requestor.

Operations that apply to a job are checked against the READ and DELETE protection specified for the queue and the owner UIC of the job. In general, READ access to a job allows you to see the attributes of the job, and DELETE access lets you delete the job.

Operations that apply to queues are checked against the WRITE and EXECUTE protection specified for the queue and the owner UIC of the queue. A user with WRITE access to a queue can submit jobs to that queue. Users with EXECUTE access to a queue may act as the operator for that queue with the ability to affect any jobs on that queue. Users with the operator (OPER) privilege have EXECUTE access to all queues. OPER privilege grants EXECUTE access to all queues. OPER privilege also enables users to establish queues and affect accounting.

• Establishing and Changing Volume Protection

VMS determines the UIC-based protection when a volume is mounted. Protection can either be defaulted from the protection recorded on the volume, or it can be explicitly specified. To change protection on a disk volume, use the SET VOLUME command. Volume protection on a magnetic volume works in a very different way. Here the protection applies equally to all files on the volume. Remember that VMS applies only READ and WRITE restrictions with respect to magnetic tapes. EXECUTE and DELETE access are meaningless. Also, users in the SYSTEM and OWNER categories are always given both READ and WRITE access, regardless of the specifications of the protection code.

If protection is not explicitly specified when a volume is initialized, all users have READ and WRITE access. If you give WRITE access to the GROUP or WORLD categories, READ access is also allowed. For magnetic tapes mounted with the FOREIGN qualifier, users in the SYSTEM and OWNER categories are always given logical and physical I/O access in addition to READ and WRITE access, regardless of what you specify in the protection code.

Note that you can change file protection on a magnetic tape only if you reinitialize the tape.

• Establishing and Changing Directory Protection

UIC-based directory file protection pertains only to disk directories and is normally established when the directory is created. At that time, the creator can either specify a protection code with the `/PROTECTION` qualifier to the DCL command `CREATE/DIRECTORY`, or he can permit the protection to default to that of the next higher directory in the tree. If the directory is a top-level directory, the protection is taken from the *master file directory (MFD)*.

Any user with `CONTROL` access can change the protection on the directory with the DCL command `SET PROTECTION`.

• Establishing and Changing UIC-Based Protection

Because UIC-based protection plays such a prominent role in controlling access to objects on your system, we'll focus for a moment on some relevant facts about how the protection is established and the ways in which it can be changed. Note that the details vary with the object to be protected.

When you create a new file, it obtains a UIC-based protection code derived from the default protection provided by the directory where it resides or from the default protection of your process.

When you create a new version of an existing file, the new file receives the protection code of the previous version of the file. You can specify a protection code when you create a copy of a file.

UIC-based protection on *global sections*, except those backed by disk files, must be reestablished every time the system is booted. If the global section is backed by a disk file, the section protection is derived from the disk file so that changing the file protection changes the section protection.

For `PSN` and page file global sections, you set the protection in the `$CRMPSC` system service call that creates the section; you cannot change the protection after the section is created.

UIC-based protection on *devices* and *logical name tables* must also be reestablished every time the system is booted. Note that you cannot change the protection on an existing logical name table.

Setting UIC-based protection on a *queue* is accomplished with the `/PROTECTION` qualifier to the `INITIALIZE/QUEUE`, `START/QUEUE`, or `SET QUEUE` command.

• Access Control Lists (ACLs)

Access control lists offer an alternative method of file protection. ACLs operate in conjunction with UIC-based protection to restrict access in very specific ways. By using ACLs, you can precisely match the specific access you want to grant or deny to specific users for each object.

To understand how this works, it is necessary to take a look at the rights database. The *rights database* is a file associating users with special names, called *identifiers*, which they are allowed to hold.

An identifier may represent a user's username and UIC, or it may represent a more general name held by many users. The latter type of identifier promotes flexibility, because it allows grouping of users according to the specific use they must make of objects on the system. Users requiring access to certain types of data would probably hold the same identifier. Because one user can be a member of several groups, users can easily access and share information.

The rights database is maintained by the system manager, who adds and removes identifiers in response to changing circumstances.

There are three types of identifiers:

1. *UIC identifiers*, which depend on the user identification codes (UICs) that uniquely identify each user on the system.
2. *General identifiers*, which are defined by the security manager in the system rights database to identify groups of users on the system.
3. *System-defined identifiers*, which describe certain types of users based on their use of the system.

When you log in, the identifiers you hold in the rights database are copied into a rights list that is part of your process. VMS uses the rights list to perform all protection checks. Additional identifiers may be added to your rights list either by the VMS login software or by software specific to your installation.

The security manager decides what kinds of access to specific objects should be granted to holders of each identifier. Often many identifiers are attached to an object, so the system manager creates a list with multiple entries. Each entry defines the group of access rights to be granted or denied the holders of the identifier named in that entry. The list of entries is the ACL, and each entry is known as an *access control list entry (ACE)*.

ACLs may be created by the system by default, by the security manager for specific objects, and by users to protect their own files. However, a user cannot create or change an ACL unless he owns the associated object, or can obtain the same access as the owner.

To summarize, an ACL consists of ACEs that grant or deny specific types of access to such system objects as files, directories, or devices. The primary value of ACLs is that they protect objects in a far more flexible way than does UIC-based protection.

When a user requests access to an object protected by an ACL, the system scans each entry in the ACL from beginning to end. If a match is found, appropriate access is granted.

Note:

The system stops searching at the first match, which means that a match occurring further down the line has no effect. It is crucial, therefore, that ACEs identifying specific users should appear before ACEs identifying groups.

The type of access needed determines the type of ACE used in a given situation. There are three types of ACEs:

1. An *identifier ACE* controls the type of access granted to a particular user or group of users.
2. A *default protection ACE* defines the default protection for a directory so that the protection can be propagated to the files and subdirectories created in that directory.
3. A *security alarm ACE* provides an alarm message when an object is accessed in a designated way.

An identifier ACE may be UIC-based, general, or system-defined. The first field in an identifier ACE consists of the keyword IDENTIFIER followed by up to 60 identifiers.

Note that the six system-defined identifiers (BATCH, NETWORK, INTERACTIVE, LOCAL, DIALUP, and REMOTE) are mutually exclusive and should not be used in combination with each other. However, you can combine them with other identifiers, that is, UICs and general identifiers. Multiple identifiers are connected with plus signs (+).

The system takes the access action specified in the ACE only for the user who holds all the identifiers specified. For example, if you wanted to grant READ access to user JONES running a batch job, the ACE would be constructed as follows:

```
(IDENTIFIER+JONES+BATCH,ACCESS+READ)
```

A number of users often share the same general identifier. These users do not have to be in the same UIC-based group. Also, a single user may be associated with a number of different general identifiers as defined in the rights database. The fact that users can hold numerous identifiers affords tremendous flexibility in selecting sets of users and defining access rights for them.

The options field in an identifier ACE controls whether the ACE is propagated, can be displayed, or can be deleted. This field starts with the keyword **OPTIONS**, followed by one or more of these keywords: **DEFAULT**, **PROTECTED**, **NONPROPAGATE**, **NONE**.

The third field in an identifier ACE specifies the type of access allowed the user(s) specified in the first field. This field begins with the keyword **ACCESS** followed by a string of access types connected with plus signs. Types of access are **READ**, **WRITE**, **EXECUTE**, **DELETE**, **CONTROL**, and **NONE**.

An example of a group-defined identifier ACE is:

```
(IDENTIFIER=PERSONNEL,ACCESS=READ+WRITE+EXECUTE+DELETE)
```

An example of a system-defined identifier ACE is:

```
(IDENTIFIER=NETWORK,ACCESS=NONE)
```

An example of a UIC-based identifier ACE is:

```
(IDENTIFIER=[SALES,JONES],ACCESS=READ)
```

The order in which ACEs appear in an ACL is extremely important, because the ACL is scanned from beginning to end, and access is granted at the first match between an identifier held by the user and an identifier specified in an ACE. For example, if the three sample ACEs above comprised an ACL for a file, the file could be accessed over a network by users holding the **PERSONNEL** identifier. Also note that if the user with the UIC **[SALES,JONES]** happened to hold the **PERSONNEL** identifier, he would have **READ**, **WRITE**, **EXECUTE**, and **DELETE** access, as well as **NETWORK** access.

To avoid inadvertently granting access to a specific individual to whom you want to deny access, UIC-based identifier ACEs should always appear before general identifier ACEs in the ACL.

The **OPTIONS=DEFAULT** option of an identifier ACE allows users to define one or more default ACEs for inclusion in the ACLs for files created in a particular directory. A default ACE applies only to new files to be created, and not to existing files. For example, if you want all files in the directory **[MALCOLM]** to have an ACE granting **READ** and **WRITE** access to users with the **PERSONNEL** identifier, you could include the following ACE in the ACL for the file **MALCOLM.DIR**:

```
(IDENTIFIER=PERSONNEL,OPTIONS=DEFAULT,ACCESS=READ+WRITE)
```

This will result in all new files created in the **[MALCOLM]** directory having the following ACE:

```
(IDENTIFIER=PERSONNEL,ACCESS=READ+WRITE)
```

Note that the **DEFAULT** option doesn't appear in the file's ACE. However, any subdirectory created in the **MALCOLM** directory has the **DEFAULT** option as part of its ACE. Thus, the default ACE can be propagated throughout the entire directory tree.

Identifier ACEs can be constructed for objects other than files and directories. Specify the object type with the **OBJECT_TYPE** qualifier on the **SET ACL** command line.

The security alarm ACE allows you to request that a security alarm message be sent to the security operator's terminal if a certain type of access takes place. The action VMS takes depends on whether the alarms have been enabled through the DCL command **SET AUDIT**. If alarms have not been enabled by the security manager (who possesses the **SECURITY** privilege), they will not reach the security operator's terminal. Thus any user-initiated alarm measures must be coordinated with the security manager. Because keeping alarms enabled uses system resources, the security manager must evaluate this need on a case by case basis.

Remember that if alarms are not enabled by the security manager with the **SET AUDIT** command, they will have no effect. The format of a security alarm ACE is:

```
(ALARM,JOURNAL=SECURITY,options),accesses
```

Also note that a security alarm ACE must include either **SUCCESS** or **FAILURE** or both. For example, to request an alarm for successful **WRITE** access to a file, you would specify the following ACE:

```
(ALARM,JOURNAL=SECURITY,WRITE=SUCCESS)
```

Managing Access Control Lists

In constructing ACLs, observe the following cautions:

1. Do not assume that specifying **ACCESS=NONE** for an identifier will ensure that no user can access the object. Users in the **SYSTEM** or **OWNER** categories may be entitled to whatever access is granted by the UIC-based protection for those categories. Also, access may be granted through privileges held by users.
2. Order your ACEs with extreme caution. Remember that access depends entirely on the first match between a user's identifier and that specified in an ACL.

3. Don't place ACLs on everything. Remember that an overabundance of ACLs has some performance implications for the system. ACLs should be used primarily when the protection they offer is especially desirable.
4. Use general identifiers to create groups of users. This practice avoids unnecessarily long ACLs.
5. Update ACLs when users leave, thus maintaining the shortest and most current ACLs possible.

• Disk Scavenging Countermeasures

Disk scavenging is reading information from "deleted" files. Remember that when a file is deleted, its header is removed from the directory, but the contents remain intact until eventually overwritten. Even when overwritten, faint residual magnetic impressions remain. These impressions may be scanned and read by programs or equipment designed for this purpose. Simply deleting a file does not put it out of reach of a potential intruder.

The first line of defense against disk scavenging is controlling access to the disks. Physical security, however, is not enough, especially considering that many breaches are inside jobs.

The usual approach to the problem involves use of *erase patterns* to overwrite deleted disk space. At some sites *erase patterns* are automatically applied whenever files are deleted or purged through controls applied to the volume. It is possible, however, that you will be asked to use the */ERASE* qualifier when invoking the DCL commands SET FILE/, DELETE, and/or PURGE for selected files.

A variation of the *erase pattern* technique is *highwater marking*, which prevents users from trailing beyond the extent of the file space where they have been permitted to write. This prevents users from scavenging portions of the disk for information that they did not themselves write.

The VMS implementation of this strategy, known as *erase-on-allocate*, is enabled as the normal default at volume initialization.

• Managing Your Files for Optimum Security

Proper file management is an important aspect of file protection. The following guidelines should help you manage your files in the most secure possible manner:

- Avoid giving your files and directories names that might attract the attention of a potential intruder, for example, SECRETWEAP.MEM.
- Purge your files regularly and delete unnecessary files.
- Use the DCI command DIRECTORY/SECURITY to monitor the ownership, protection code, and ACLs on your files.
- Ensure that the protection on your mail files makes them accessible only to you and the system.
- Periodically check the dates your files were last revised and compare with your recollection of your most recent activity.
- When placing ACLs on your files, be sure you know which users hold the identifiers you have specified.
- Devote special attention to protecting files containing command procedures and executable programs.



Chapter 5 • Network Security

• Security for a DECnet Node

Networks present a more complex security challenge than does the single-system environment, primarily because of increased operational complexity and decentralization of control. This chapter provides information on how system managers can implement various features of VMS to improve the security in their networks.

• The Reference Monitor in a Network

The *reference monitor concept*, which has been discussed in the context of the single system environment, also applies to a network of interconnected computer systems. In the case of a network, a subject on one system holds a subject seeking access, while another holds the object of the access attempt.

Both source and target computer systems must have their own implementation of the reference monitor concept if there is to be security in the network. However, individual reference monitors can deal only with local subjects and objects.

Therefore, accessing an object across a network requires a phantom object on the system with the real subject, and a corresponding phantom subject on the system with the real object.

There are three critical requirements for achieving security in the network environment:

- The real subject on the source machine and the phantom object on the target machine must correspond. This correspondence must be managed by the two reference monitors and must be consistent with the security policy intended on the target machine.
- The authorization database on the target machine must express an access authorization for a phantom object that corresponds to the correct real subject on the correct source machine.
- A protected means of communications between the two reference monitors must exist to establish reliable correspondence between real and phantom subjects.

• Establishing Subject Correspondence

VMS and DECnet-VAX provide several mechanisms for establishing a correspondence between a subject or process on a source node and another on a target node.

The default account mechanisms allow any subject on any node to be placed in correspondence with a default subject on a target node. This subject can, in turn, gain access to objects on behalf of a requesting subject and return the required information.

Note that in this mechanism any subject can be placed in correspondence with a default subject on a target node. There is little selectivity in the establishment of the correspondence.

At the other extreme is the use of explicit or username/password access control on the establishment of a subject at the target node. This mechanism restricts access to those objects accessible to the named user, but it has the undesirable effect of broadcasting passwords over the network.

• Proxy Accounts

Use of proxy accounts is another option for establishing correspondence between subjects. This option requires the target reference monitor to maintain a table of source subjects and the corresponding local user names. Each request from a subject on a source node is mapped into the creation of a subject representing the corresponding target user. This offers the explicit control associated with username/password control while adequately protecting passwords.

With proxy logins, there is no need to embed passwords in commands to copy a file. Also, a file's protection code need not be set to allow the WORLD category of users READ access to transfer a file. Instead, the user simply issues the following form of the DCL command COPY.

```
COPY remote-node::file-spec file-spec
```

Setting Up Proxy Logins

Two utilities are used to set up proxy logins: AUTHORIZE and NCP (Network Control Program). You may want to create a command procedure to assist you in implementing proxy access through a step-by-step approach.

The command procedure could provide the following functions:

1. Check if the proxy is turned on for your system.
2. Create a proxy account for sharing files.
3. Add a user to access a proxy account.
4. Remove a user from access to a proxy account.

4. Remove a user from access to a proxy account.
5. List users authorized to access a proxy account.

To set up proxy logins without using a command procedure, use `AUTHORIZE` to create or modify the network proxy authorization file, `NETPROXY.LIST`, that contains the names of all users allowed proxy access to the system and the names of all proxy accounts defined for the remote users.




Figure 5-1 User Account Routing

5.4 • Network Security

Set up a proxy account on your node for use by one or more users at other nodes as follows:

1. Decide the purpose of the account, the name of the local account, and which foreign users will be admitted.
2. If the local account doesn't exist, create it with `AUTHORIZE`. If the account does exist, examine it to ensure that it is adequately restricted. Proxy accounts should be restricted so that they prohibit interactive users and batch jobs, which is to say they should permit only network logins.
3. Review the privileges on the account. Generally, you should avoid granting privileges to proxy login accounts.
4. If the network user authorization file does not already exist, create it with the `AUTHORIZE` command `CREATE/PROXY`.
5. Add as many network user authorization records as necessary with the `AUTHORIZE` command `ADD/PROXY`.




Figure 5.2 Proxy Account Routing

6. Check the default protection on the directory and customize it as necessary.
7. Examine any command procedure used at login time and specified by /LGICMD. The command procedure should reside in a well-protected directory owned by a user other than the owner of the proxy account. It should also prohibit WRITE access for those who use the account.
8. Inform the remote node's security manager of the identities of users from that node who have been authorized for access to your node.

Remember that AUTHORIZE performs certain automatic maintenance functions on the NETUAF.DAT proxy login file. Whenever the username changes through a RENAME or COPY command, the associated change is made in the NETUAF. Similarly, when you remove an account from SYSUAF.DAT, all entries for which there is a matching local username are removed from NETUAF.DAT.

Proxy access is in effect a selective merging of the authorization databases of the affected systems. Therefore, overall security is only as good as that of the least secure node involved.

• Specifying Authorizations

The approach used to specify authorizations for access to objects depends on the mechanism used for establishing correspondence between subjects. The various default account mechanisms create anonymous subjects on the target node. As a result, objects that are to be made accessible to a default account must permit the WORLD user category full access, thus leaving the object unprotected.

If explicit access control or proxy access is used to establish correspondence between subjects, the authorization can be granted to the target subject selected by the username or proxy. In this case, the full range of VMS authorization mechanisms can be used.

• Protecting Communications

The security of network operations partly depends on the ability of source and target reference monitor mechanisms to communicate in a secure manner. An intruder must not be able to observe passwords or mimic a source node that has been granted proxy access.

Encryption of communications adds a significant measure of security to network operations.

The *Enhanced Ethernet Security System*, composed of DESNC security network controllers and the VAX/KDC Ethernet security manager software, provides security management and control in an Ethernet LAN.

Enhanced Ethernet Security System

DESNC controllers for Ethernet local area networks upgrade security through authentication of Ethernet nodes, enforcement of a mandatory access policy among Ethernet nodes, and data protection through encryption and integrity controls. DESNCs provide transparent cryptographic security at a level consistent with the National Institute of Standards and Technology (NIST) Data Encryption Standard (DES).

With the Enhanced Ethernet Security System, you can logically separate different classes of nodes in a LAN. For example, development systems and production systems can exist physically on the same Ethernet but are not allowed to send Ethernet packets to each other. PCs and workstations can be similarly managed with the Enhanced Ethernet Security System.

The DESNC controller is a store-and-forward device providing real-time cryptographic processing of Ethernet frames (messages) over an Ethernet LAN (local area network). Decryption restores the data to its original form through a client node that takes advantage of encryption/decryption services.

A DESNC controller resides between the client node and a backbone Ethernet. It intercepts data sent from the node, verifies the identity of the node, and encrypts the data before sending it to the network. The DESNC controller also intercepts data received from the network, decrypting it before passing it to the client node requiring protection.

A DESNC controller can accommodate any combination of workstations, servers, or VAX processors.

Encryption is implemented at the Data Link Layer of the International Standards model. Thus, protocols residing above Ethernet may be used separately or simultaneously.

Operating system software that can run on client nodes supported by the DESNC controller includes VMS, RSX, and ULTRIX, as well as industry standard operating systems such as UNIX* and MS-DOS.†

VAX/KDC is a layered product that serves as the central authority for managing the DESNC controllers and enforcing security policy for an Ethernet LAN or extended LAN. The function of VAX/KDC is to perform management and distribution services for a group of DESNC controllers. Multiple VAX/KDC nodes may be used to improve the availability of the networked DESNC controllers.

* UNIX is a registered trademark of the American Telephone and Telegraph Company.

† MS-DOS is a registered trademark of the Microsoft Corporation.




Figure 5.3 Security Controls for Ethernet




Figure 5-4 ISO Reference Model for Open Systems Interconnection (OSI)

VAX/KDC software contains four critical database capabilities:

- *Access control policy*, which establishes Data Link level connection privileges for each node connected to a DESNC controller
- *Node information*, which includes all configuration data concerning nodes attached to each DESNC in the network
- *Status*, which indicates the current status of each DESNC in the network
- *Audit trail*, which provides a history of all relevant security events under the control of the system manager

The Security Modem

Network security can also be enhanced by the use of modems especially designed to prevent unauthorized access. Digital's *Scholar Plus* modem features an access security system which protects data with a system of callback and password verification. Unauthorized access is prevented by means of modem parameters, dial memory, and callback memory locations.

Mandatory password access prevents unauthorized modification when the modem is unattended. The callback security feature screens incoming calls before allowing access to the host system, and provides an audit trail of access attempts. Flexibility is facilitated by four security levels:

- Level 1 allows pass thru by means of a password.
- Level 2 allows callback to a predefined phone number with password security.
- Level 3 allows callback to a predefined phone number with password security and phone number validation assigned to the password.
- Level 4 allows callback to a predefined or user-defined phone number with password security.

The *Scholar Plus* modem enforces security with little inconvenience to an authorized remote user, who simply inputs the correct password and/or telephone number for validation by the modem. Upon validation, the *Scholar Plus* performs a callback or pass thru, and the user is on-line and ready to transmit data. The *Scholar Plus* automatically adjusts to the operating speed of Digital's terminals for all speeds between 1200 and 9600 bits per second.

• Security for Packet-Switched Data Networks

A fundamental technology used in data networks is *packet switching*. With this technology, user data and the accompanying control information needed to ensure delivery are formed in discrete entities called *packets*. The network dynamically interweaves the packets of many users over shared transmission facilities and routes the packets to their destinations.

Packet-switched data networks (PSDNs) can be public or private. When connected to the PSDN, calls may be connected to any other number on that network, or to any number on another public PSDN that is cross-connected to it. Calls may be normally or reverse-charged. The customer pays a tariff related to time-of-day, number called, duration, and amount of data.

Private PSDNs are owned and operated by one company or group of companies, and are used for intercompany rather than intracompany communications. Otherwise, they offer the same general features as their public counterparts.

The *VAX/Packetnet System Interface (VAX/PSI)* serves as an interface to the worldwide X.25 networks. *X.25 networks* are public data networks that provide a data connection between different types of computers by using the CCITT suite of protocols. Primary applications are terminal connections and electronic mail. VAX/PSI provides a Security utility and an Accounting utility.

System managers, using the security system, can allow calls into PSI in two ways:

- To other data terminating equipments (where a DTE is a computer or terminal connected to the Packetnet system) that make calls to VAX/PSI (These are incoming calls.)
- To users of the host system who make calls to VAX/PSI (These are outgoing calls.)

The security system has no effect until it is turned on by data entering the security databases. When the security system is turned on, there is total surveillance. Calls are not allowed on the VAX/PSI software unless specifically permitted.

These three databases are used by the PSI security system:

- The *Agent Rights Database*, which lists agents and their rights concerning VAX/PSI (An agent is either a host user or a remote DTE.)
- The *DTE Access Control Database*, which contains a list of remote DTEs plus an access control list for each DTE.
- The *Destination Access Control Database*, which contains a list of PSI destinations or applications, plus an access control list for each application.

Access control entries (ACEs) comprise the access control lists (ACLs) that protect particular system objects. ACLs grant or deny these types of access:

- NONE, for no access to PSI
- INCOMING, for incoming non-reverse-charge calls
- OUTGOING, for outgoing reverse-charge calls
- INCOMING+REVERSE-CHARGE, for all incoming calls including reverse-charge calls
- OUTGOING+CHARGE, for all outgoing calls

The system searches each entry in the ACL in the usual manner, that is, from first to last, for the first match it can find. Access is granted or denied on the basis of the first match encountered.

The PSI Accounting utility can be used to record details of how PSI is used. You may use this feature to provide an audit trail, to establish records for the purpose of charging users for X.25 resources, or to compile performance data.

The Accounting utility can record data that allows you to calculate the cost of any call and to determine who was using the network at any given time.

You may also record all calls (whether completed or failed) and all access to permanent virtual circuits (PVCs) on multihop, native mode, and PSI Access VAX/PSI systems.

• Terminal Servers

Local area networks (LANs) allow computing resources to be physically distributed throughout a facility. A typical LAN can have 1000 or more attachments on a single coaxial cable over one mile long. A potential problem in such installations is the limited bandwidth of the LAN. For this reason communication architectures operating in this shared environment should use the bandwidth efficiently. The *Local Area Transport (LAT)* architecture, used by Digital in many of its Ethernet products, has been designed to satisfy this goal.

A *terminal server* contains specialized software that performs a dedicated function. LAT terminal servers provide the facilities of a network terminal switch.

Groups are subdivisions of a LAT network. Groups do not affect the physical makeup of the network, but they can be used to logically partition the network into combinations of services, service nodes, and server ports. Groups define the access that server port users and service nodes have to the network. Each service node and the services it offers are in one or more groups, and each terminal server port is in one or more groups.

If a server port and a service share the same group, then the following is possible:

- The port user can display information about the service node and its services.
- The port user can connect to a service offered by the service node.
- The service can send a host-initiated request to the port.

The network manager may assign a unique service group for each service on the LAN and assign that service group to all service nodes offering that service. Alternatively, if a network has a large number of services, it may be better to assign groups to sets of services rather than to individual services. This is especially useful when several nodes offer the same service, such as a VAXcluster offered as a single service.

Effective assignment of authorized groups requires detailed information about LAT services on the LAN and about each port user.

If you want to restrict a terminal to a single LAT session with one service, specify a dedicated service for the port. Use a dedicated service when you want to simulate a hard-wired connection between a terminal and a service node. Users that require access to only one service should be assigned to a port that is dedicated to that service.

Two important security features of the DECserver system are security status and passwords. The server manager controls security by assigning each user to a particular status level and by creating passwords.

Security status levels restrict the server commands available to the port user. One of three levels of security apply to each port:

- Secure status (most restrictive)
- Nonprivileged status (less restrictive)
- Privileged status (no restrictions)

Choosing the security level of server ports is an important aspect of managing the server. The SECURITY port characteristic allows you to set up the security levels on a port-by-port basis. With this feature, you can limit the use of server commands, particularly the privileged SET commands that change the operational database. Portions of the DECserver command set are associated with each of the three security levels. Thus there are privileged commands, nonprivileged commands, and secure commands.

- *Privileged commands* include the entire command set.
- *Nonprivileged commands* are a subset of the privileged commands.
- *Secure commands* are a subset of the nonprivileged commands.

Typically, this is how you decide on privilege levels for users:

- The server manager can use all server commands (the privileged command set).
- Most users at interactive terminals on DECserver ports can use the nonprivileged commands.
- Users that you want to restrict can use the secure commands only. By default, all server ports are nonprivileged, and only the server manager should modify the security status of a port.

Privileged status lets a port accept all server commands: privileged, nonprivileged, and secure. Privileged status is the only level allowing access to the SET and DEFINE commands that manage the server and all its ports. Privileged commands let you do everything you can do with nonprivileged commands plus the following:

- Configure ports.
- Establish security levels for ports.
- Customize the permanent and operational databases.
- Perform tests.
- Observe the status of the server, its ports, and the LAN network environment.

A port with *nonprivileged status* has access to all the nonprivileged commands required for accessing and using services, and to a few SET PORT and DEFINE PORT commands. Nonprivileged status also allows access to some SHOW commands. In addition, this status permits users to use some restricted server features such as broadcasting to other ports. Nonprivileged commands let a user do the following:

- Establish characteristics of his port.
- Obtain a display of available services.
- Establish sessions with services.
- Switch between sessions.
- Display many characteristics of his port and its sessions.

Secure status restricts commands available on a port to a subset of the nonprivileged commands. This subset contains commands required for accessing and using services from a particular port, and for specifying some characteristics of that port. Secure users can use SHOW PORT only for their own port and they can change only a few port characteristics. Secure status is useful for isolating port users from some features of the server and from other users.

Like security status options, password options are important to your management of the server. The DECserver system uses the following types of password:

- Privileged passwords
- Login passwords
- Lock passwords

Passwords help you control the server's use and protect the server, its operational database, and the users' efforts. The commands that set passwords are all privileged commands except for that which sets a lock password, which is a general user tool.

The *privileged password* prevents misuse of the server management commands. You can change this password in both the operational and management databases. Any user who knows the privileged password can issue the SET PRIVILEGED command at any interactive terminal to turn his port into a privileged port. (More than one port can be privileged at any time.) Because users can enter commands intended for server management at privileged ports, the privileged password must be well protected.

The *login password* prevents unauthorized use of the server. If you use the login password for a port, any potential user trying to log in is prompted for the login password. The server will complete the login only if the correct password is supplied. The login password is especially valuable in restricting access to terminals located in a public place. The login password can be changed in both the operational and permanent databases.

A single login password is used for the whole server, although the password is enabled on a port-by-port basis. If you plan to enable the login password at one or more ports, the password should be well protected. In addition, you should change the password regularly and inform affected users of the new password.

Any user can specify a *lock password* to prevent unauthorized access to his port. When entering the LOCK command, the user types and then verifies a lock password. The server does not accept input from the user's keyboard until the user again enters the same password to unlock the terminal.

• DECnet-VAX Accounts

DECnet-VAX accounts permit certain types of access to your system from remote nodes without requiring account and password information. Instead, this information is specified in the DECnet-VAX executor and object databases. These accounts are controlled through the system authorization file using techniques similar to those used for captive user accounts.

When setting up accounts for DECnet-VAX use, it is advisable to follow these guidelines:

1. DECnet-VAX has no requirement for a privileged default account, and you should not provide one. In addition, create a default account for objects only when required.
2. UICs of the network nonprivileged accounts should be unique for each group and user. Also, the group code must exceed the system UIC group number to avoid granting the SYSTEM user category for file access to the user.
3. Keep the privileges for DECnet-VAX accounts to a minimum. Typically, this means giving only TMP MBX (Temporary Mailbox) and NET MBX (Network Mailbox) to unprivileged accounts.
4. Maintain the secrecy of passwords for the DECnet-VAX accounts.
5. The account for the FAL (File Access Listener) object should have a group code in its UIC that differs from that of every other account in the system, including accounts for other DECnet objects.
6. The member number of the owner UIC of the default directory for the FAL account should differ from that of the owner UIC of the FAL account.

The DECnet-VAX Database

The DECnet-VAX node and circuit databases control how other computers are allowed to connect to your computer.

Because a computer connection permits unwanted assaults on both your own security and that of any other computer in the network, it requires very strict control.

To promote the security of the databases, observe the following guidelines:

1. Define receive and transmit passwords for all nodes in the database. When possible, the transmit and receive passwords should be different and not obvious.
2. Always enable verification on any circuit that goes outside a locked computer room, or goes to a machine with a different security environment.
3. Do not define default access rights in the database for external nodes.
4. In general, do not enable backup synchronous dial-up for autoanswer. Systems that have incoming dial-up for production purposes should control which nodes can connect.

• Network Usage

Network use is restricted in many countries, either by law or by contract with the major communications division of the government. For example, several countries have laws to protect personal data and impose restrictions on moving such data across national boundaries. Security managers should be familiar with all applicable laws and ensure that their sites are in compliance.

• Security Concerns on a Cluster

Clustered VMS systems use VMS hardware and software to share disks, resources, and a common operating system among various VAX computers. In this case, the computers are said to be joined in a VAXcluster.

There are two types of VAXcluster: homogeneous and nonhomogeneous (or heterogeneous). A *homogeneous VAXcluster* has an identical operating system environment on each member node. A *nonhomogeneous VAXcluster* has a unique environment on each node.

From a security standpoint, the node's part in a VAXcluster has little significance, because each node operates as a single system. All security features previously described, therefore, also apply to any node in the cluster. The only difference is that when a security manager implements any feature on one node of a homogeneous cluster, all others are affected. The reason for this is that each cluster node mediates access by its subjects to all objects in the cluster.

In effect, the cluster operates within a single security perimeter. The reference monitor on each node acts as a gateway through that perimeter.

There is one area, however, in which the actions the cluster manager takes in setting up the VAXcluster can affect the security operations of the system. This area is the creation and management of the various elements of the overall authorization database.

Management of security on a VAXcluster is facilitated by ensuring that all elements of the user authorization data exist in a common database.

Authorization elements include the system and network user authorization files and the rights database, which are present on all VMS systems, and the optional autologin file.

If you create the optional autologin file, consider maintaining it in a common authorization database along with your authorization files and rights database.

Also remember that on a clustered system the autologin file must include the cluster node name as a prefix to the terminal name.

• Using DECnet Between Cluster Nodes

Although VAXclusters offer communication facilities for the most common operations, such as file sharing and lock management, other VMS features may require DECnet for use across a cluster. Examples of intrachuster use include the need to access disks not cluster-accessible and higher level features available through DCL commands such as SHOW USERS.

For maximum coherence in DECnet operations, set up a proxy database that maps users into their own accounts when they initiate DECnet operations. For each node in a homogeneous cluster, you would add a proxy file record using the AUTHORIZE command:

```
ADD/PROXY node:: *
```

If you are running a nonhomogeneous cluster, a more complex arrangement of proxies is needed to cross-map your users only as they are authorized.

• Summary

Security operations are enhanced on a VAXcluster when all the authorization data resides on a common shared disk.

Remember that each user must have the same UIC, group number, and set of identifiers defined on each cluster node.

On a shared disk, the protection of a file from a specific user cannot effectively exceed the maximum access that user can gain from one of the nodes.

VMS security features operate on clustered systems just as they do on nonclustered systems.

Chapter 6 • Auditing

• Introduction

The auditing features of VMS can combat security threats associated with user irresponsibility and probing, and also create and maintain records of system activity. Some events, such as uses of privilege and certain login failures, are always auditable; others can be specified by users or system managers.

Intelligent auditing means striking a balance between the desire to record everything and the need to restrict alarms to a number that can be investigated effectively. The best policy may be to implement alarms for only a few events at any given time.

• Auditing with Security Alarms

Security alarms are messages sent to the security operator terminal when specific events take place.

Alarms can be used to detect intrusion attempts by unauthorized users, to monitor undesirable activity at a site, and to conduct general auditing of various aspects of system use. For example, an alarm might be enabled that sends a message to the security operator terminal whenever a change is made to a UAF record.

Effective implementation of VMS auditing features requires that events to be audited be carefully selected, that the appropriate alarms be enabled, that a security operator terminal be enabled, and that alarm information be intelligently utilized.

Events that can be audited are:

- Selected types of access to files and global sections
- An event requested by an ACL on a file or global section
- Use of privilege to access files and global sections
- Installation of images
- Logins, logouts, and break-in attempts
- Modifications to the system and network UAF
- Changes to system and user passwords
- Execution of the SET AUDIT command
- Volume mounts and dismounts

The information included in the message depends on the type of event. However, all alarm messages contain the following elements:

- **OPCODE** heading, which includes the date and time the alarm was sent
- **Type of alarm**
- **Date and time the alarm occurred**
- **Perpetrator of the event**, as identified by the username and process identification (PID)

The following are some suggestions for enabling security alarms:

- **Enable security auditing for login failures (LOGFAIL) and multiple failures**, which are identified as break-in attempts (BREAKIN). This is the best way to detect probing by both outsiders and insiders.
- **Enable security auditing for LOGIN**. Auditing successful logins, especially those from REMOTE and DIALUP sources, helps track which accounts are being used. Note that an audit record is written before a user who logs into a privileged account can disguise his identity.
- **Enable the FILE=FAILURE type of security audit**. This technique audits all file protection violations and is an excellent way to catch probes.
- **Apply ACL-based file access auditing to detect WRITE access to critical system files**. Also consider auditing successful access to detect penetrations or failed access to detect probing.
- **Audit accesses to files containing especially critical data**.
- **Audit use of privilege to access files**. Be advised, however, that this class of auditing may generate a large volume of output, because privileges are routinely used in normal system operation for such tasks as mail delivery and operator backups.

• Audit Reduction Facility

As previously stated, information resulting from enabled alarms is written by the operating system into the operator's log file, which typically contains a large volume of data. The audit reduction facility lets you extract only the specific kind of information you want at any given time. For example, you can extract all security alarm records generated by user Jones after November 1. Another command might extract all security alarm records generated by break-in attempts, any access to a file using the SYSPRV privilege, or any access to a file using the BYPASS privilege.

It is also possible to audit an entire terminal session. Auditing of terminal sessions for selected users can be enforced by using a special captive account and appropriate command procedures. A user for whom auditing is enforced must first log into the captive account and then into his own account. The captive account ensures that the session is audited.

In addition to security alarms, VMS provides other data which may be useful in tracking system activity. The system accounting log contains records of all system job terminations, including all INTERACTIVE, BATCH, and NETWORK jobs, as well as prior jobs and other process terminations. Optionally, activations of all or selected images may be recorded in the accounting log.

Most network operations, such as mail delivery and access to files from remote nodes, initiate a network server job for which a log file is created. This log file is normally named NETSERVER.LOG and is located in the default directory of the account under which the job ran. NETSERVER.LOG may be useful in tracking events that were initiated over the network.

Chapter 7 • Security for the User

• Introduction

From the point of view of the system, there are two kinds of people in the world: authorized users and unauthorized users.

Unauthorized users have no business using the system, and the security features of VMS are intended to keep them out. As an *authorized user*, it is your responsibility to use the system in a secure manner. This section of the handbook explains security features relevant to your use of the system.

• Logging In

Initial access to the system is achieved by logging in. At login time, you must establish that you are an authorized user. Usually you accomplish this by providing a username and a password.

Login also allows the system to impose restrictions. The procedure may be relatively simple at a minimum security site, or highly complex in a maximum security environment. There are seven types of login. Typical user logins are LOCAL, DIALUP, and REMOTE. Logins usually performed by the system are NETWORK, BATCH, DETACHED, and SUBPROCESS. Logins may be interactive or noninteractive.

An *interactive login* requires you to provide information in response to system prompts. For example, if the words "SELECT SYSTEM" appear on the terminal, you type the name of the system. You might then type your username at the "Username:" prompt, and your password at the "Password:" prompt. If the system is satisfied with the information you provide, it grants you whatever access you are authorized to have.

A *noninteractive login* is performed by the system without user interaction.

Classes and types of login are summarized as follows.

LOCAL	Interactive
DIALUP	Interactive
REMOTE	Interactive
NETWORK	Noninteractive
BATCH	Noninteractive
DETACHED	Dependent on parent process
SUBPROCESS	Noninteractive

We will now look at some of the elements of a typical interactive login.

7.2 • Security for the User

You may first see an announcement message identifying the node and, if relevant, the cluster that you have successfully accessed. The announcement message immediately precedes the "Username:" prompt.

After you have logged in you may see a welcome message, which often states the software version of VMS you are using, and possibly the name of the node as well.

After the announcement message, you may see information relating to the last successful login. This information might include the time of the last LOCAL, DIALUP, or REMOTE login, the time of the last successful noninteractive login, and the number of login failures. Login failures in this context are those caused by providing incorrect passwords.

Your system manager may choose to suppress announcement and welcome messages in order to avoid revealing information to a potential intruder. This is particularly likely at medium to high security sites.

Although last login and failure messages may also be suppressed, their display promotes good security because they keep users informed of access attempts. Also, by demonstrating that the system is monitored, they may discourage potential intruders.

• Passwords

Good password security plays an essential role in protecting the system from unauthorized access. You undoubtedly have noticed that your password does not appear on your terminal screen when you type it in. This prevents an onlooker from seeing your password.

Various kinds of passwords may be required to access the system. Most systems require a *user password*, which is your personal password. In addition, you may need to provide a *system password* to log in to a particular terminal. In certain high security situations, the system may demand both a *primary password* and a *secondary password* instead of one user password.

Passwords are stored by the system in an encrypted form. This prevents an intruder from accessing a list of actual passwords. Because the encryption is one-way, an intruder cannot deduce actual passwords by analyzing their encoded equivalents.

System passwords control access to particular terminals. Although they may be used simply to tighten security in a general sense, they are most often required to limit access to terminals which may be special targets for unauthorized use.

The four types of account available on VMS have distinctive password requirements. The four types of account are:

1. *Open accounts*, which require no password
2. *Captive accounts*, which permit very limited operations and may require a password
3. Accounts that require a password but prohibit a user from changing it
4. Accounts that require passwords that the user or system manager can change (the most common type of account, and the type this handbook assumes you have)

When an account is opened for you, you are usually assigned a password. You should change this password the first time you log in.

Unless your system manager requires use of the random password generator, you will choose your own password. In choosing a password, avoid any word which may be associated with you, such as your nickname, the name of your pet, the make of car you drive, and so on.

It is advisable to change your password from time to time. You can accomplish this with the DCL command SET PASSWORD. If the policy in effect at your site allows you to choose your own password, remember to observe the minimum character length requirement.

If your system demands use of the automatic password generator, you may choose from among several potential new passwords that are generated for you. If you don't like any of the initial choices presented, you can have additional password choices generated.

The password generator produces character strings that resemble English words, but cannot be found in a dictionary because they are not real words. This thwarts potential intruders who may attempt access by using small computers programmed to enter every word in the dictionary as a password.

A password may include digits as well as letters. Including digits in your password makes password guessing much more difficult.

As stated above, some terminals require both a primary and secondary password before granting access. In some cases, one user will enter both passwords, thereby enhancing security for the obvious reason that two passwords are harder for an intruder to guess than one. Most often, however, two people, each of whom has been entrusted with one of the passwords, will access the system. This involves visual contact between the users, thereby minimizing the possibility of an imposter effecting unauthorized access. This procedure is relatively cumbersome, and is used only in high security environments, or to protect exceptionally sensitive accounts.

You should never reveal your password to anyone. Passwords are more often given away than stolen. Remember that an intruder can use his access to your account to severely damage the system. It is a serious mistake to believe that your password simply grants access to your personal account. Bear in mind that once you've revealed your password, you have lost all control over its distribution. Maintaining the secrecy of your password is your responsibility.

If You Have Multiple Accounts on Different Systems

Most of the comments above assume you have only one account. However, you may hold more than one account on different systems. In this case, observe the following guidelines:

1. If any one of the systems demands especially high security, use a unique password for that system.
2. If any of the systems uses non-Digital computers, use unique passwords on your accounts.
3. If all the systems use VMS or have equally good password encryption mechanisms, you can use the same password for all accounts provided you maintain good password security.
4. If any system does not have a password encryption mechanism, use a unique password on that system.

PASSWORD ENCRYPTION




Figure 7-1 Password Encryption

Summary of Password Guidelines

1. Select passwords that cannot be easily guessed.
2. Never write your password down.
3. Never tell anyone your password.
4. Do not include your password on any file, including electronic mail messages.
5. Change your password frequently.
6. Change your password immediately if you think it may have been discovered.
7. Avoid using the same password on multiple systems unless you have good reason to believe that doing so is safe.

Causes of Login Failures

For a variety of reasons, you may occasionally experience difficulty logging in. Some common login failure symptoms and their causes are detailed in the following table:

Symptom	Cause
No response from the terminal	Attempting to use a defective terminal
No response from any terminal	Attempting to log in when the system is down
Message: User authorization failure	Mistyping the username or password
Message: User authorization failure	Attempting to use an expired account
Message: User authorization failure	Attempting to use an expired password

Other possible reasons for login failure are the need for a system password, shift restrictions, prohibitions on certain types of login, or system activation of evasion measures.

System Password Failures

Your login attempt will fail if the terminal you are trying to use requires a system password. The system will not tell you that a system password is needed, so it may appear to you that the system is down. If you suspect this may be the problem, and if you have not been given the system password, try another terminal.

If the system password you have been given doesn't seem to work, the password may have been changed. In this case, try another terminal that doesn't demand a system password, or take steps to find out what the new password is.

Login Class Restrictions

You may not be able to log in because you are restricted from certain types of logins. As you recall, there are five basic types of logins: LOCAL, REMOTE, DIAL-UP, BATCH, and NETWORK. The general term INTERACTIVE may be used to include or exclude LOCAL, REMOTE, and DIAL-UP access. If restrictions are in effect, you may, for example, be unable to access the system over a network because the system manager has denied you this form of access.

Shift Restrictions

Your system manager can restrict your logins to certain hours per day and/or days per week. If this is the case, your login attempts outside those periods will fail. The restrictions may affect all classes of logins without reservation, or various types of logins may be restricted in different ways. If this is the problem with your login attempt, you will be so informed by the system.

Please remember that when this type of restriction applies to batch jobs, any job you submit to be run outside your authorized hours will not run. Note also that the job will not be automatically resubmitted to run during authorized hours. Similarly, any job you submit which runs into an unauthorized period will be immediately aborted at the end of the authorized period. To avoid problems, be sure that you are aware of any restrictions that may affect your use of the system.

Dial-Up Login Failures

Your security manager can control the number of dial-up login failures permitted before the connection is broken. If your login fails and you have some tries remaining, just press the return key and try again. If only one failure is permitted, or you've exhausted the limit, you must redial and start again.

Limiting dial-up failures before breaking the connection discourages would-be intruders from guessing a password by entering many possible choices in succession.

Note that this feature cannot prevent this practice, because the intruder can always re-dial and start again. Like many security measures, the idea is to make things tougher on the unauthorized user.

Break-In Evasion

VMS provides evasion measures in the event of repeated login failures. If you can't log in, it's possible that someone has made numerous attempts to log in using your username in conjunction with invalid passwords. Of course, the problem might simply be your many typos while trying to enter your password. The evasion mechanism responds by denying access for a period of time, so that you can't log in even with the correct password. Report this situation to your security manager so that he can investigate the possibility that illicit access has been attempted.

A summary of the possible symptoms and causes of login failures just discussed follows:

Symptoms and Causes of Login Failures

Symptom	Cause
No response from terminal.	The device requires a system password.
No response from terminal after entry of system password.	The system password changed and you were not notified.
Message: Not authorized to log in from this source	The attempted class of login (LOCAL, DIAL-UP, REMOTE, BATCH, or NETWORK) is prohibited.
Message: Not authorized to log in at this time.	The day of the week or hours of the day are not permitted for you for this class of login.
Message: User authorization failure (and no known user failure occurred).	An apparent break-in has been attempted at the terminal using your username, and the system has temporarily disabled all logins at that terminal by your username.

• Network Security Considerations for Users

Networks present special concerns in controlling access to data. The challenge is somewhat similar to that of securing a single-system environment, but is more demanding because of increased operational complexity and decentralization of control. In addition, technological limitations make data sent across a network vulnerable to interception.

There are, however, ways to tighten network security. As a user, you should know about:

- Access control strings in file specifications and command procedures
- Proxy login
- Proper use of the VMS Mail Utility

Network access control strings are designed to be included in the file specifications of DCL commands that work over the DECnet-VAX network.

Access control strings permit a user on a local node to request an operation using a file on a remote node. The string consists of the remote node name, the username for the remote account, and the user's password. Because they contain enough information to allow anyone to break into the remote account, they create a significant security exposure.

You should take all possible steps to protect access control strings. For example, avoid leaving information in hardcopy or on video terminals. You should also avoid placing networking commands in command procedures where they would be targets for discovery.

Remember that the syntax that requires the username and password to be placed in quotation marks and followed by two colons attracts the attention of a password-hunting intruder. If you must place access control strings in a command procedure, implement good file protection.

A more secure approach might be to explore with your system manager the possibility of using a proxy login account instead of access control strings. *Proxy logins* permit access without revealing the information contained in network access control strings, that is, usernames and passwords.

Before a user can issue a request that initiates a proxy login, the system manager must create a proxy account for the user. This involves creation of a *network user authorization file (NETUAF)* that identifies remote users who should be granted access to the proxy account.

The main advantage of a proxy login is that it avoids forwarding the password as a string.

Setting up proxy accounts requires time and effort on the part of the system manager, but the expenditure pays off in greatly enhanced network security. However, it is inadvisable to implement proxy logins for privileged accounts.

Mail files are choice targets for security assaults. With that in mind, use discretion in the contents of your messages. For example, never include your password or details about using your account. It is usually a good idea to delete your messages after you read them. This is especially important when the messages contain sensitive information. Take particular care in disposing of hardcopies. If you need to retain some mail files, implement good file protection.

• Logging Off the System

Never leave your terminal unattended without logging off. Doing so exposes the entire system to unauthorized access. Remember that one such lapse may render futile all other security measures enforced at your site.

An issue to consider is how much information can be safely left on your terminal screen after you have logged out. In a minimum security environment, this may not be a major consideration. In a medium to high security environment, it is a major consideration.

If your site falls into the medium security category, you are well-advised to leave nothing but the logout message on your screen.

If you work in a high security environment, you may be directed to turn off your terminal every time you log out. This practice deprives a potential intruder of any useful information, such as a valid username, that would otherwise be visible on the terminal.

The preceding suggestions involve effort on your part. We will therefore focus for a moment on why security is important to you, as an individual and as part of a group. One of the many destructive things an intruder can do is to delete or damage files. This could represent a very personal loss if the files belong to you.

Of course, the havoc created by wholesale damage to the system may, at the least, cause serious inconvenience to yourself and your coworkers. But a security breach can do a lot more than cause inconvenience. Your company or organization may be so badly hurt by disclosure of secret data or disruption of operations that it may be unable to function effectively. For example, theft of proprietary information could have a devastating effect on your firm's position in its industry. This might adversely impact your salary and benefits, and could even cost you your job.

Another consideration is that if users fail to voluntarily observe good security practices, management has little choice but to impose additional mandatory protection mechanisms. This makes using the system more tedious for all users.

You can help prevent security-related difficulties by following the guidelines recommended at your site, and by implementing the suggestions described in this handbook.

Appendix A • Running VMS in a C2 Environment

Many of the security features of VMS were specifically designed to meet the requirements for a class C2 system as defined in the *Department of Defense Trusted Computer System Evaluation Criteria*, published by the National Computer Security Center. VMS was formally evaluated at the C2 level by the NCSC in August of 1986.

This appendix describes how the features of VMS relate to the C2 security model, and notes considerations for operating a VMS system within the C2 framework.

• The Trusted Computing Base

The Trusted Computing Base (TCB) provided by VMS encompasses much of the operating system, including:

- The entire executive and file system
- All other system components that operate in inner access modes
- Most system programs installed with privilege
- A variety of other utilities used by system managers to maintain data relevant to the TCB

The objects for which VMS provides full C2 protection are files and directories that are accessible through normal file access techniques or through global sections, as well as devices, logical name tables, and queues. VMS also fulfills all other requirements demanded by the the C2 criteria, including auditing of accesses.

• Protecting the TCB

The code and data that make up the VMS TCB reside in files and, in part, in the address space of the running operating system. Integrity of code and data is protected by file access controls and memory page protection. Memory page protection is set up by VMS as it executes and is normally not of concern to the system manager. The files containing the TCB are correctly protected by default when the system is installed. This protection can, however, be changed by sufficiently privileged users.

Certain privileges allow their holders to bypass normal file and memory access controls either directly or indirectly. These privileges, therefore, should be granted only to the system manager, the system security officer, or other highly trusted persons.

Privileges in the FILES and ALL categories allow the holder to violate the integrity of the TCB.

Privileges in the SYSTEM category allow the holder to interfere with normal system operation and cause denial of service. They do not allow the user to directly violate object access controls. However, some irregular uses of privileges in the SYSTEM category may, by very indirect means, ultimately result in violating of access controls.

Privileges in the DEVOUR and GROUP categories permit the user to consume resources without limits, which may deny service and interfere with the operations of others in the same group. In particular, the GRPPRV privilege allows a user to violate normal controls within his group.

• Individual Accountability

VMS enforces accountability by means of usernames, UICs, and passwords. The following practices and features, however, result in loss of individual accountability and must not be used in a C2 environment:

- The same UIC assigned to more than one user
- Open accounts
- Group accounts
- Autologin
- Network proxy accounts for groups

• Object Protection and Reuse

Reuse of system memory pages is protected by the memory management subsystem and cannot be defeated. Reuse of disk blocks is protected by the highwater marking and erase on delete features. Conformance with C2 criteria is achieved by enabling of highwater marking, which is the default.

VMS views magnetic tapes as single-user devices. Tape protection is available only at the volume level. Therefore an entire volume may be assigned ownership and protection, but not the individual files contained therein. As a consequence, VMS provides no protection against reuse of tape. Tapes that are recycled to new users must be externally erased by operations personnel.

• Protection of the Audit Trail

The security audit trail is recorded in the operator log file and on terminals enabled as security operators.

The operator log is normally protected against reading or modification by unauthorized users. It is possible to protect the contents of the audit log with the following measures:

1. Place a hardcopy terminal enabled as a security operator in a physically secure location.
2. Protect the audit log file with the following audit measures:
 - Enable audits on ACL and audit events with the following command:
`$ SET AUDIT /ALARM /ENABLE=(ACL,AUDIT)`
 - Place on the operator log file an ACL entry that enables auditing of all accesses for modification and deletion.

These steps ensure that attempts to tamper with the audit log result in the system audit controls being left obviously turned off, or with a last "foot-print" in the audit log. Circumventing these measures requires extensive programming.

• Auditing Actions of a System Operator or Administrator

All actions taken by such trusted users as operators, administrators, and security officers can be audited by enforced use of terminal session auditing. Attempts to defeat the auditing can be detected by measures similar to those used to protect the audit log:

1. Enable auditing of authorization modifications.
2. Place ACL entries on the captive login command procedures and the directories containing them to detect modification of the procedures.

• Documentation

The *Trusted Facility Manual* consists of the *Guide to VMS System Security* and the applicable reference documentation. The first four chapters of the *Guide to VMS System Security* constitute the *Security Features User's Guide* and should be available to all system users.

• Physical Security

Physical and environmental security are critical to the secure operation of the system. Serious attention should be given to preventing theft of media and output. In addition, the console terminal must be physically secured because it controls the operation of the CPU and, consequently, the operation of the system.

• Configuration Guidelines

The security features described in this handbook apply to most VAX configurations. They are supported by all VAX CPUs, including MicroVAX CPUs, and apply to all supported mass storage and communications devices. VAXcluster configurations also fully support the security features. A VAXcluster is considered a single security and management domain and normally operates with a shared authorization database.

The NCSC evaluation criteria do not address network operation. However, when connected to a DECnet network, VMS provides security commensurate with the security of the base operating system if the following restrictions are met:

- All operating systems connected to the network are VMS systems or systems of equivalent security and are systems administered in a secure manner.
- Default accounts are not provided for file and general task access. Limiting access to explicit access control strings and to proxy access preserves individual accountability.
- The communications lines are secured from wiretaps by use of link encryption devices or by physical security.

Appendix B • Managing Classified Data

The discretionary access controls and other security features of VMS are geared to the needs of most commercial and government installations. They allow the system manager to implement the elements of his security policy in an appropriate, effective, and flexible manner.

Computer systems entrusted with classified data, however, require additional security safeguards to meet the various levels of protection established by the Department of Defense.

VMS has been formally evaluated at the C2 level as defined in the *Department of Defense Trusted Computer System Evaluation Criteria* published by the National Computer Security Center. Security enhancements to VMS fulfill most of the criteria of the B1 level with reference to mandatory access controls, labeling, and auditing.

Because sophisticated security features have been engineered into VMS, such enhancements are accomplished primarily with refinements to existing capabilities, rather than by means of an "add-on" security package. In addition, latent capabilities in the operating system are developed by replacing or adding VMS system components to achieve certain objectives, such as labeled object protection.

The VMS Security Enhancement Service (VMS/SES) is a software security consulting package that provides many features of mandatory access controls and security auditing for the VMS operating system. VMS/SES combines the services of a trained Digital consultant with licensed software and documentation. It provides a system administrator or system security officer with the means to devise a systemwide security policy and help safeguard users, data, and software from security threats.

VMS/SES is intended for government agencies, national defense organizations, and prime contractors who need to label and protect classified information processed on VMS systems.

VMS/SES also provides an effective means of evaluating the effects of mandatory access controls on application design and system management.

The primary features of VMS/SES are:

- Technical support provided by a Digital consultant
- Custom installation of VMS/SES software

B.2 • Managing Classified Data

- Enhancement of operating software to help resist unauthorized access to data
- Addition of security labels to classified data by a nondiscretionary control package
- Provision of labeling capability and mandatory access controls to help enforce user authorization boundaries
- Enhanced administrative control through:
 - Checking user clearance before granting access to data
 - Restricting output to authorized devices
 - Labeling printed output with correct classification
 - Auditing file activity

As a component of the service, Digital's VMS/SES consultant reviews system security policies and controls to assess existing system security and demonstrate the mapping of security requirements onto the capabilities of VMS and VMS/SES. The consultant then installs the application software and trains the system manager and users in optimal use of the software.

Digital's VMS/SES consultant provides support in the following areas:

- System security review assesses the state of existing system security. The results of the review are presented in a written report in a standardized format.
- System security planning to help map the general security requirements of the system onto VMS/SES. The consultant prepares a written report in standardized format that describes how these requirements can be met by VMS Security Enhancements Service.
- User orientation provides an overview of the user interface and operational characteristics of the VMS/SES software.
- Security and system manager orientation provides a thorough grounding in management and operations of the VMS/SES software.
- Licensed software installation consists of software installation and technical guidance.

VMS/SES helps enforce security policy through a series of validation and access checking mechanisms. It also contains a flexible secure print facility and enhanced security auditing.

VMS/SES is designed for secure data processing environments that require mandatory access controls. It helps the system security officer to define and control access between subjects and objects, and to designate sensitivity labels for any user, file, or device that consists of hierarchical levels and nonhierarchical categories or compartments.

With VMS/SES, security level and category names can be stored in the VMS rights database, and security ranges for users can be assigned by an AUTHORIZE command. LOGINOUT then enables users to log into any classification within their authorized range.

Users can be divided into two groups: nonprivileged and privileged. The VMS security reference monitor helps to ensure that nonprivileged users cannot violate the mandatory access control policies it implements.

Privileged users have elevated privileges that permit them to change the operation of the system. A management guide is provided to help establish system configuration and to provide privileged users with system setup guidelines.

The following example illustrates how files are protected: When a file is created, it inherits the user's security label. When a volume is initialized, it can be labeled with a minimum and maximum security label. From that point on, only files that fall into the volume's initialized security range can be created on the volume. When a subject requests access to an object, mandatory access control mechanisms in the software compare labels and grant or deny access through a protection algorithm.

Special features and commands of VMS/SES software include:

- Dynamic SYSGEN parameter permits enabling or disabling of the mandatory access mediation mechanism.
- Labeled object protection uses the Digital Command Language (DCL) to label and protect objects and devices.
- Classification labels support 256 security levels and up to 64 integrity and security categories. (Classification labels optionally support 128 security categories if integrity categories are not used.)
- The DOWNGRADE privilege allows a process to write to objects at a lower security level than itself, or to reclassify an object to a lower security level. The BYPASS privilege can be used to circumvent mandatory access controls.
- The AUTHORIZE command supports the classification of users with levels and categories, and provides a means to enter class labels into the rights database.
- BACKLIP saves and restores security labels of files and volumes.
- The SECURITY ENHANCED PRINT FACILITY prints security classifications on hardcopy output. It allows for security header pages, page banners, and inclusive page numbering.
- AEDIT enhancements allow monitoring of files for mandatory access. Such auditing can be enabled for a given level or range of levels.

VMS/SES software is supported on VAX and MicroVAX computer systems.



Glossary

Access control list:

A list that defines the kind of access to be granted or denied to users of an object. Access control lists can be created for objects such as files, devices, and mailboxes. Each access control list consists of one or more entries known as access control list entries.

Access control list entry:

An entry in an access control list. Access control list entries may specify identifiers and the access rights to be granted or denied the holders of the identifiers, default protection for directories, or security alarm details. Access control lists for each object can hold many entries, limited only by overall space and performance considerations.

ACE:

See access control list entry.

ACL:

See access control list.

ACL Editor:

A VMS utility that helps users create and maintain access control lists.

Alarm:

See Security Alarm.

Alphanumeric UIC:

A format of user identification code (UIC) that specifies the user's group and member number in alphanumeric form rather than numeric form.

Attribute:

In the security context, an attribute is a field of information maintained in the rights database that identifies some characteristic accorded to all holders of the identifier. For example, if an identifier possesses the resource attribute, holders of that identifier can change resources such as disk space usage to that identifier.

Auditing:

The act of noting the occurrence of an event that has security implications.

Authentication:

The act of establishing the identity of users when they start to use the system. VMS (and most other commercial operating systems) use passwords as the primary authentication mechanism.

Breach:

A break in the system security that results in admittance of a person or program to an object.

Break-in Attempt:

An effort made by an unauthorized source to gain access to the system. Since the first system access is achieved through logging in, break-in attempts primarily refer to attempts to log in illegally. The attempts focus on supplying passwords for users known to have accounts on the system, through informed guesses or other trial-and-error methods.

Captive Account:

A type of VMS account that limits the activities of the user. Typically, the user is restricted to using certain command procedures and commands. The user may not be allowed to use the CTRL/Y key. (This type of account is synonymous with a turnkey or tied account.)

Decryption:

The process that restores encoded information to its original unencoded form.

Discretionary Controls:

Security controls that are applied at the user's option; that is, they are not required. Access control lists are typical of such optional security features. Discretionary controls are the opposite of mandatory controls.

Disk Scavenging:

A term that refers to any method of obtaining information that the owner intended to discard from a disk. The information, although no longer accessible to the original owner by normal means, can be retrieved and used by one of the scavenging methods because a sufficient amount of its original magnetic encoding remains.

Encryption:

A process of encoding information so that its content is not immediately obvious to anyone who obtains a copy of it.

Erase-on-allocate:

A technique that applies an erasure pattern when a new area is allocated for a file's extent. The new area is erased with the erasure pattern so that subsequent attempts to read the area yield only the erasure pattern and not some valuable remaining data. This technique is used to discourage disk scavenging.

Erase-on-delete:

A technique that applies an erasure pattern when a file is deleted or purged. This technique is used to discourage disk scavenging.

Erasure Pattern:

A character string that can be used to overwrite magnetic media to erase the information previously stored in that area.

Evasive Action:

A responsive behavior by VMS to discourage break-in attempts when they appear to be in progress. VMS has a set of criteria it uses to detect the fact that break-in attempts may be underway. Typically, when VMS becomes suspicious that an unauthorized user is attempting to log in, the evasive action locks out all login attempts by the offender for a limited period of time.

General Identifier:

One of three possible types of identifiers that specify one or more groups of users. The general identifier is alphanumeric and typically is a convenient term that symbolizes the nature of the group of users. For example, typical general identifiers might be PAYROLL for all users allowed to run payroll applications or RESERVATIONS for operators at the reservations desk.

Highwater Marking:

A technique for discouraging disk scavenging. This technique tracks the furthest extent that the owner of a file has written into the file's allocated area. It then prohibits any attempts at reading beyond the written area, on the premise that any information that exists beyond the currently written limit is information that some user had intended to discard. VMS accomplishes the goals of highwater marking with its erase-on-allocate strategy.

Holder:

A user who possesses a particular identifier. The term holder is used in conjunction with the term identifier. Users are said to be holders of identifiers if they possess the identifiers. The rights database is the place in the system where the associations of users and the identifiers they hold are permanently kept. However, each process also has a rights list that includes all the identifiers the process is authorized to hold.

G.4 - Glossary

Identifier:

A notation that defines a user or group of users. There are three types of identifier: UIC identifiers, system-defined identifiers, and general identifiers.

Locked Password:

A password that cannot be changed by the account's owner. Only system managers or users with the SYSPRV privilege can change locked passwords.

Login:

The series of actions involved in authenticating a user to the system and creating a process that runs on the user's behalf.

Mandatory Controls:

Security controls that are imposed by the system on all users. Mandatory controls are the opposite of discretionary controls. Mandatory controls are implemented in the VMS Security Enhancement Service.

Non discretionary Controls:

See Mandatory Controls

Open Accounts:

Accounts that do not require passwords.

Passwords:

Character strings that users provide at login time to validate their identity and prove their authorization to access the account. There are system passwords and user passwords. User passwords include both primary and secondary passwords.

Primary Password:

The first user password requested from the user. Systems may also require a secondary password. This password must be associated with the user name that is supplied with it.

Privileges:

A means of protecting the use of certain system functions that can affect system resources and integrity. System managers grant privileges according to the user's needs and deny them to restrict the user's access to the system.

Proxy Login:

A login type that permits a user from a remote node to log in to a local node as if the user owned an account on the local node. However, the user does not specify a password in the access control string. The remote user may own the account or share the account with other users.

Rights Database:

The collection of data the system maintains and uses to associate identifiers and the holders of the identifiers with their rights and attributes.

Rights List:

The list associated with each process that includes all the identifiers the process holds.

Secondary Password:

A user password that may be required at login time immediately after the primary password has been correctly submitted. Primary and secondary passwords supplied by separate users can ensure that more than one user is present at the login. A less common use requires a secondary password to increase the password length so that the greater number of combinations makes password guessing more time-consuming.

Secure Terminal Path:

VMS software designed to ensure that users can log in only to terminals that are already logged out. When the user presses the BREAK key on a terminal, the secure terminal path (if enabled) responds by first disconnecting any logged-in process and then initiating a login. If a process is not logged in at the terminal, the login can proceed immediately. Note that the secure terminal path can be used only with terminals that are connected directly to the system.

Security Alarm:

A message sent to operator terminals that are enabled as security operators. Security alarms are triggered by the occurrence of an event previously designated as worthy of the alarm because of its security implications.

Security Manager:

In the VMS context, the person or persons responsible for protecting the security of the computer system. This role is sometimes performed by the same person who functions as a system manager. It requires the same skills as the system manager, but also includes additional privilege (the SECURITY privilege) and knowledge of the security features provided with the VMS operating system.

Security Operator Terminal:

A class of terminal that has been enabled to receive messages sent by OPICOM to "security operators". These messages are security alarm messages. Normally such a terminal is a hardcopy terminal in a protected room. The output provides a log of security-related events and details that identify the source of the event.

G-6 • Glossary

System-Defined Identifier:

One of three classes of identifiers. System-defined identifiers are provided by the system to identify groups of users according to their usage of the system. For example, all users who access the system by dial-up receive the DIALUP identifier.

System Password:

A password required by a terminal before login can be initiated.

Tied Account:

See Captive Account

Trojan Horse Program:

A program that gains access to otherwise secured areas through a pretext of serving one purpose when its real intent is devious and potentially damaging

Turnkey Account:

See Captive Account

UIC:

See User Identification Code

User Identification Code:

A coded notation that represents a user of the system. Normally, each user has a unique identification code, although at a few sites some users may share the same UIC. In that case, the system cannot distinguish one user from another. User identification codes include a designation of the user and the user's group.

User Password:

A password that is associated with a user. This password must be correctly supplied when the user attempts to log in so that the user is authenticated for access to the system. The two types of user passwords are known as primary and secondary, which also designate the sequence in which they are entered.

Glossary - G-7

Virus:

A form of Trojan horse program that replicates itself from user to user and system to system. See Trojan Horse Program.

Worm:

A command procedure or executable image written and placed on the system for the sole purpose of seeking unauthorized access to files and accounts on the system. The "worm" seeks access to a user file through a flaw in the file protection. If successful, the worm modifies the file so that it carries a copy of the worm. Each time an unsuspecting user executes the code that contains the worm, the worm attempts to propagate itself into the other poorly protected procedures or images, travelling along a path known as a worm-hole. The worm seeks a procedure that runs from a privileged account so that the worm can inflict damage to the system security.

Bibliography

DATAPRO Reports on Information Security, DATAPRO Research, Delran, NJ 08075.

DECnet/E V4.0 System Manager's Guide, Digital Equipment Corporation, Part Number AA-H503C-TC.

Department of Defense Trusted Computer System Evaluation Criteria, DoD 5200.28-STD, December, 1985.

Guide to VMS System Security, Digital Equipment Corporation, Part Number AA-L840A-TE.

Rdb System Manager's Guide, Digital Equipment Corporation.

VAX ALL-IN-1 System Manager's Guide, Digital Equipment Corporation, Part Number AA-Y187B-TE.

VAX Key Distribution Center Security Manager's Guide, Digital Equipment Corporation, Part Number AA-KM69A-TE.

