

AD-A111 578

FORD AEROSPACE AND COMMUNICATIONS CORP PALO ALTO CA W--ETC F/8 9/2
K505 COMPUTER PROGRAM DEVELOPMENT SPECIFICATIONS (TYPE B-S). (K--ETC(U)
DEC 80


MDA903-77-C-0533

UNCLASSIFIED


NDL-TR7932-REV-3

ML

For
Order



11157



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

ADA 111 578

2

SECURE MINICOMPUTER OPERATING SYSTEM (KSOS)
COMPUTER PROGRAM DEVELOPMENT
SPECIFICATIONS (TYPE B-5)


Department of Defense Kernelized Secure Operating System

- I. SECURITY KERNEL (CDRL 0002AF)
- II. UNIX EMULATOR (CDRL 0002AG)
- III. SECURITY-RELATED SOFTWARE (CDRL 0002AH)

Contract MDA 903-77-C-0333

Prepared for:

Defense Supply Service—Washington
Room 1D245, The Pentagon
Washington, D.C. 20310



DTIC FILE COPY

Approved for public release; distribution unlimited.


Ford Aerospace &
Communications Corporation
Western Development
Laboratories Division

3939 Fabian Way
Palo Alto, California 94303

SECURE MINICOMPUTER OPERATING SYSTEM (KSOS) SECURITY KERNEL COMPUTER PROGRAM DEVELOPMENT SPECIFICATION (TYPE B5)

Department of Defense Kernelized Secure Operating System

Contract MDA 903-77-C-0333
CDRL 0002AF

Prepared for:

Defense Supply Service - Washington
Room 1D245, The Pentagon
Washington, DC 20310

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
For	
Dist. by/	
AV. Codes	
or	
Dist. by/	
Dist. by/	
A	



**Ford Aerospace &
Communications Corporation**
Western Development
Laboratories Division

3939 Fabian Way
Palo Alto, California 94303

NOTICE

The Department of Defense Kernelized Secure Operating System (KSOS) is being produced under contract to the U.S. Government. KSOS is intended to be compatible with the Western Electric Company's UNIX[®] Operating System (a proprietary product). KSOS is not part of the UNIX license software and use of KSOS is independent of any UNIX license agreement. Use of KSOS does not authorize use of UNIX in the absence of an appropriate licensing agreement.

This document is a major revision of an earlier draft dated July 1978. That and other previous editions are obsolete and should not be used.

UNIX and PWB/UNIX are trade/service marks of the Bell System.

DEC and PDP are registered trademarks of the Digital Equipment Corporation, Maynard MA.

TABLE OF CONTENTS

1.	SCOPE	1
1.1	Identification	1
1.2	Functional Summary	1
1.3	Terminology	1
2.	APPLICABLE DOCUMENTS	5
2.1	Government Documents	5
2.1.1.	Directives, Manuals and Standards	5
2.1.2.	Reports	5
2.2	Non-Government Documents	6
2.3.	Other References Not Part of This Specification	6
3.	REQUIREMENTS	8
3.1.	Computer Program Definition	8
3.1.1.	Interface Requirements	8
3.1.1.1.	Interface Block Diagram	8
3.1.1.2.	Detailed Interface Definition	8
3.1.1.2.1.	Kernel Objects	9
3.1.1.2.1.1.	Kernel Object: Processes	10
3.1.1.2.1.2.	Kernel Object: Process Segments	12
3.1.1.2.1.3.	Kernel Object: Files	14
3.1.1.2.1.4.	Kernel Object: Devices	15
3.1.1.2.1.5.	Kernel Object: File Subtypes	15
3.1.1.2.2.	Process Structure	16
3.1.1.2.3.	Hardware Interface	17
3.1.1.2.4.	Emulator/User Mode Software Interface	17
3.1.1.2.5.	Privileged Subsystem Interfaces	18
3.1.1.2.6.	Secure Terminal Interface	20
3.2.	Detailed Functional Requirements	21
3.2.1.	Kernel function: K_build-segment	21
3.2.1.1.	Inputs	21
3.2.1.2.	Processing	22
3.2.1.3.	Outputs	23
3.2.2.	Kernel function: K_close	23
3.2.2.1.	Inputs	23
3.2.2.2.	Processing	23
3.2.2.3.	Outputs	24
3.2.3.	Kernel function: K_create	24
3.2.3.1.	Inputs	24
3.2.3.2.	Processing	24
3.2.3.3.	Outputs	24
3.2.4.	Kernel function: K_device_function	25
3.2.4.1.	Inputs	25
3.2.4.2.	Processing	25
3.2.4.3.	Outputs	26
3.2.5.	Kernel function: K_fork	26
3.2.5.1.	Inputs	26

3.2.5.2.	Processing	26
3.2.5.3.	Outputs	27
3.2.6.	Kernel function: K_get_file_status	27
3.2.6.1.	Inputs	27
3.2.6.2.	Processing	27
3.2.6.3.	Outputs	28
3.2.7.	Kernel function: K_get_object_level	28
3.2.7.1.	Inputs	28
3.2.7.2.	Processing	28
3.2.7.3.	Outputs	28
3.2.8.	Kernel function: K_get_process_status	29
3.2.8.1.	Inputs	29
3.2.8.2.	Processing	29
3.2.8.3.	Outputs	29
3.2.9.	Kernel function: K_get_segment_status	30
3.2.9.1.	Inputs	30
3.2.9.2.	Processing	30
3.2.9.3.	Outputs	30
3.2.10.	Kernel function: K_interrupt_return	31
3.2.10.1.	Inputs	31
3.2.10.2.	Processing	31
3.2.10.3.	Outputs	31
3.2.11.	Kernel function: K_invoke	31
3.2.11.1.	Inputs	31
3.2.11.2.	Processing	32
3.2.11.3.	Process Bootstrapper Processing	32
3.2.11.4.	Outputs	33
3.2.12.	Kernel function: K_link	33
3.2.12.1.	Inputs	33
3.2.12.2.	Processing	33
3.2.12.3.	Outputs	34
3.2.13.	Kernel function: K_mount	34
3.2.13.1.	Inputs	34
3.2.13.2.	Processing	34
3.2.13.3.	Outputs	36
3.2.14.	Kernel Function: K_nap	36
3.2.14.1.	Inputs	36
3.2.14.2.	Processing	36
3.2.14.3.	Outputs	36
3.2.15.	Kernel function: K_open	37
3.2.15.1.	Inputs	37
3.2.15.2.	Processing	37
3.2.15.3.	Outputs	38
3.2.16.	Kernel function: K_post	38
3.2.16.1.	Inputs	38
3.2.16.2.	Processing	39
3.2.16.3.	Outputs	39
3.2.17.	Kernel function: K_read_block	39
3.2.17.1.	Inputs	39
3.2.17.2.	Processing	39
3.2.17.3.	Outputs	40
3.2.18.	Kernel function: K_receive	41
3.2.18.1.	Inputs	41
3.2.18.2.	Processing	41
3.2.18.3.	Outputs	41
3.2.19.	Kernel function: K_release_process	41

3.2.19.1. Inputs	41
3.2.19.2. Processing	41
3.2.19.3. Outputs	42
3.2.20. Kernel function: K_release_segment	42
3.2.20.1. Inputs	42
3.2.20.2. Processing	42
3.2.20.3. Outputs	42
3.2.21. Kernel function: K_remap	42
3.2.21.1. Inputs	42
3.2.21.2. Processing	43
3.2.21.3. Outputs	43
3.2.22. Kernel function: K_rendezvous_segment	44
3.2.22.1. Inputs	44
3.2.22.2. Processing	44
3.2.22.3. Outputs	45
3.2.23. Kernel function: K_secure_terminal_lock	45
3.2.23.1. Inputs	45
3.2.23.2. Processing	45
3.2.23.3. Outputs	46
3.2.24. Kernel function: K_set_file_status	46
3.2.24.1. Inputs	46
3.2.24.2. Processing	47
3.2.24.3. Outputs	47
3.2.25. Kernel function: K_set_object_level	47
3.2.25.1. Inputs	47
3.2.25.2. Processing	48
3.2.25.3. Outputs	48
3.2.26. Kernel function: K_set_process_status	48
3.2.26.1. Inputs	48
3.2.26.2. Processing	49
3.2.26.3. Outputs	49
3.2.27. Kernel function: K_set_segment_status	50
3.2.27.1. Inputs	50
3.2.27.2. Processing	50
3.2.27.3. Outputs	50
3.2.28. Kernel function: K_signal	51
3.2.28.1. Inputs	51
3.2.28.2. Processing	51
3.2.28.3. Outputs	51
3.2.29. Kernel function: K_spawn	51
3.2.29.1. Inputs	52
3.2.29.2. Processing	52
3.2.29.3. Outputs	52
3.2.30. Kernel function: K_special_function	52
3.2.30.1. Inputs	52
3.2.30.2. Processing	53
3.2.30.3. Outputs	53
3.2.31. Kernel function: K_unlink	53
3.2.31.1. Inputs	53
3.2.31.2. Processing	53
3.2.31.3. Outputs	54
3.2.32. Kernel function: K_unmount	54
3.2.32.1. Inputs	54
3.2.32.2. Processing	54
3.2.32.3. Outputs	55
3.2.33. Kernel function: K_walk_process_table	55

3.2.33.1. Inputs	55
3.2.33.2 Processing	55
3.2.33.3. Outputs	55
3.2.34. Kernel function: K_write_block	56
3.2.34.1. Inputs	56
3.2.34.2. Processing	56
3.2.34.3. Outputs	56
3.2.35. Kernel Resource Manager	57
3.2.36. Secure Terminal Interface	57
3.2.37. Audit Messages	57
3.2.38. Kernel Trap Management	58
3.2.38.1. Inputs	58
3.2.38.2. Processing	58
3.2.39. Pseudo Interrupts	59
3.2.40. Special Requirements	59
3.2.40.1. Human Performance	60
3.2.40.2. Government Furnished Property List	60
3.3 Adaptation	60
3.3.1. General Environment	60
3.3.2. System Parameters	60
3.3.3. System Capacities	60
4. QUALITY ASSURANCE PROVISIONS	61
4.1. Introduction	61
4.1.1. Category I Testing	61
4.1.2. Category II Testing	61
5. PREPARATION FOR DELIVERY	64
6. NOTES	65
10. APPENDIX	66
Formal Specifications	68

Kernelized Secure Operating System (KSOS):
Security Kernel Computer Program Development Specification (Type B5)

1. SCOPE

1.1 Identification

This specification establishes the performance, design, development and test requirements for the Kernelized Secure Operating System (referred to as "KSOS") Security Kernel. The KSOS Security Kernel is intended to support the emulation of the UNIX operating system while providing an environment suitable for the construction of new secure systems.

1.2 Functional Summary

The primary purpose of this specification is to define the functional requirements for a Security Kernel to support the emulation of the standard UNIX operating system. This specification further defines the design, test, performance, and quality assurance provisions for the KSOS Security Kernel.

This specification is intended to serve an additional purpose, that of defining a particular implementation of the KSOS Security Kernel on a Digital Equipment Corporation PDP-11/70 computer system. Thus, there are a number of PDP-11/70 specific points throughout the specification. For implementations of the KSOS Kernel on other hardware bases, these points should be treated as non-binding, informative commentary.

The KSOS Security Kernel shall support the emulation of the standard UNIX operating system within the constraints of the multilevel security model. The Security Kernel shall provide all the data and computational objects required to construct a general purpose operating system. The Security Kernel shall mediate all information exchanges within the system permitting only those exchanges which are consistent with the multilevel security model. The Security Kernel shall provide secure information channels which permit the user to communicate directly with "trusted" software processes. The secure information channels shall not be subject to compromise due to "spoofing" by untrusted software.

This specification is organized in the following way. First the relevant Kernel requirements are detailed. The presentation of the requirements includes the general interface requirements, interface block diagram, and the detailed interface definition. The detailed interface definition specifies the Security Kernel call interface and significant Kernel-related software elements, i.e., the scheduler and secure terminal interface. System adaptability is discussed followed by the specification of quality assurance provisions for the Security Kernel.

1.3 Terminology

The language used throughout this specification attempts to conform to the guidelines of Section 3.2.3 of MIL-STD-490. In particular, the word "shall"

Security Kernel B5 Specification

means that the specification expresses a provision that is binding. The words "should" and "may" mean that the specification expresses a provision which is non-mandatory. The word "will" is used to express a declaration of purpose on the part of the Government.

KSOS shall support the mandatory DoD security policy of DoD Directive 5200.1-R that is embodied in a Government approved mathematical model. [Verif 78] Proofs of the system's security properties shall be in terms of this model. KSOS shall provide a minimum of eight (8) hierarchical security classification categories (or simply security categories) and a minimum of thirty-two (32) mutually independent security compartments. The security categories shall be such that

UNCLASSIFIED < CONFIDENTIAL < SECRET < TOP SECRET

Where "<" is defined in accordance with the requirements of DoD 5200.1-R. One security compartment shall be reserved for read protection of system data bases and programs. This compartment shall be called the "system" compartment. KSOS shall provide for Kernel-enforced integrity. Integrity is defined as the formal mathematical dual of security [Biba 75]. At least four (4) hierarchical integrity classification categories (or simply integrity categories) shall be provided in KSOS. The integrity categories include at least the following three classification categories:

USER < OPERATOR < ADMINISTRATOR

Although there is no immediate requirement for integrity compartments, the development contractor should include provisions to ease their later inclusion. For the remainder of this specification, the term "security level" means the combination of a security category and a set of security compartments (which may be null). The term "integrity level" means the combination of an integrity category, and a (presently always null) set of integrity compartments. The term "level" (or "access level") means the security and integrity level, that is, the combination of a security category, a set of security compartments (which may be null), an integrity category, and a (presently always null) set of integrity compartments. The "access matrix" defines for a particular user or group the combination of the maximum security level permitted for that user or group, all security compartments to which that user or group has access, the maximum integrity level for that user or group, and all integrity compartments (presently always null) to which that user or group has access.

In a secure system it is necessary to have processes external to the kernel which perform security critical functions. These have come to be called "trusted processes". In KSOS there is a refinement in this terminology. As Figure 1-1 shows there are subdivisions of the software inside the security perimeter. Some of the processes are "privileged". This means that they may violate one or more of the security rules enforced by the Kernel. These processes must be designed and implemented with the same care and thoroughness as the Kernel itself. Section 3.1.1.2.5 presents the privileges which may be

Security Kernel B5 Specification




Figure 1-1. Terminology Figure

afforded processes. A second subdivision has been termed "responsible" processes. These are processes are not privileged to violate any of the Kernel's rules, but which are nonetheless important to the overall security of

Security Kernel B5 Specification

the system. Examples of responsible processes include the Secure Mail facility and the processes which manipulate the data bases describing the users' security and integrity permissions (the "access matrix"). These processes must also be carefully designed and implemented, but do not require formal specification or verification. Because they do perform security-related functions, they cannot include untrusted components (e.g. the UNIX Emulator). The third subdivision is the "encapsulated utilities" (EU in the figure). These are self contained functions to which the user simply supplies parameters. The encapsulated utilities are not privileged to violate Kernel rules, and do not affect the security of users. The chief function in the encapsulated utilities is system generation.

Security Kernel B5 Specification

2. APPLICABLE DOCUMENTS

The following documents, of exact issue shown, form a part of this specification to the extent specified herein. In the event of a conflict between the referenced documents and the contents of this specification, this specification shall be considered a superseding requirement. In the text references to these documents are in the form [Name date], e.g. [Biba 75].

2.1 Government Documents

2.1.1 Directives, Manuals and Standards

- a. DoD 5200.1-R Information Security Program Regulation
- b. DoD 5200.28 Security Requirements for Data Processing (ADP)
- c. DoD 5200.28-M ADP Security Manual
- d. MIL-STD-483 Configuration Management
- e. MIL-STD-490 Specification Practices
- f. MIL-STD-1521A Technical Reviews and Audits

2.1.2 Reports

- a. [A-Specs 78] "KSOS System Specification (Type A)", WDL-TR7808 Revision 1, Ford Aerospace and Communications Corporation, Palo Alto, CA (July 1978).
- b. [Bell and LaPadula 73] Bell, D.E. and LaPadula, L.J., "Secure Computer Systems", ESD-TR-73-278, Volume I-III, MITRE Corporation, Bedford, MA (November 1973 - June 1974).
- c. [Biba 75] Biba, K.J., "Integrity Considerations for Secure Computer Systems", MTR-3153, MITRE Corporation, Bedford, MA (June 1975).
- d. [Emulator 78] "KSOS Computer Program Development Specification (Type B5): UNIX Emulator", WDL-TR7933, Ford Aerospace and Communications Corporation, Palo Alto, CA (September 1978).
- e. [NKSR 78] "KSOS Computer Program Development Specification (Type B5): Non-Kernel Security-Related Software", WDL-TR7934, Ford Aerospace and Communications Corporation, Palo Alto, CA (September 1978).
- f. [Verif 78] "KSOS Verification Plan", WDL-TR7809, Ford Aerospace and Communications Corporation, Palo Alto, CA (March 1978).
- g. [Walter et al. 74] Walter, K.G. et al., "Primitive Models for Computer Security", ESD-TR-74-117, Case Western Reserve University, Cleveland, OH (January 1974).

Security Kernel B5 Specification

2.2 Non-Government Documents

- a. [BBN 75] "Interface Message Processor, Specifications for the Interconnection of a Host and an IMP", Report 1822, Bolt Beranek and Newman Inc., Cambridge, MA, (December 1975).
- b. [Holmgren et al. 77] Holmgren, S.F., Healy, D.C., Jones, P.B. and Kasprzycki, E., "Illinois Inter-Process Communication Facility for UNIX", CAC Technical Memorandum 84, CCTC-WAD Document 7507, Center for Advanced Computation, University of Illinois at Urbana-Champaign, Urbana, IL (April 1977).
- c. [Lampson 73] Lampson, B., "A Note on the Confinement Problem", CACM, Volume 16, Number 10, pp 613 - 615 (October 73).
- d. [Lipner 75] Lipner, S.B., "A Comment on the Confinement Problem", Proc. Fifth Symposium on Operating Systems Principles, ACM SIGOPS Review, Volume 9, Number 5, pp 192 - 196 (19-21 November 1975).
- e. [Nemeth et al. 77] Nemeth, A.G., Sunshine, C., Zucker, S. and Tepper, S., "Progress Towards a DeFacto DoD UNIX Inter-Process Communication Standard", Working Paper, (May 1977). (Available from the authors.)
- f. [Parnas 72] Parnas, D.L., "A Technique for Software Module Specification with Examples", CACM, Volume 15, Number 5, pp 330 - 336 (May 1972).
- g. [Postel 78a] Postel, J.B., "Internetwork Protocol Specification", Version 4, Information Sciences Institute, University of Southern California, Marina del Rey, CA (February 1979).
- h. [Postel 78b] Postel, J.B., "Specification of Internetwork Transmission Control Protocol - TCP", Version 4, Information Sciences Institute, University of Southern California, Marina del Rey, CA (February 1979).
- i. [Roubine and Robinson 77] Roubine, O., L. Robinson, Special Reference Manual, 3rd ed., Technical Report CSG-45, SRI International, Menlo Park, CA (January 1977).
- j. [Sunshine 77] Sunshine, C., "Interprocess Communication Extensions for the UNIX Operating System: I Design Considerations", R-2064/1-AF, RAND Corporation,
- k. [Zucker 77] Zucker, S., "Interprocess Communication Extensions for the UNIX Operating System: II Implementation", R-2064/2-AF, RAND Corporation, Santa Monica, CA (June 1977).

2.3 Other References Not Part of This Specification

- a. [Thompson 75] Thompson, K. and Ritchie, D.M., "UNIX Programmer's Manual", Sixth Edition, Western Electric Corporation, Greensboro, NC (May 1975).
- b. [Ritchie 74] Ritchie, D.M. and Thompson, K., "The UNIX Timesharing System", CACM, Volume 17, Number 5, pp 365 - 375 (May 1974).

Security Kernel B5 Specification

- c. [UNIX 75] "UNIX Program Listings", Version 6.0, Western Electric Corporation, Greensboro, NC (May 1975).

Security Kernel B5 Specification

3. REQUIREMENTS

3.1 Computer Program Definition

This Computer Program Configuration Item (CPCI) defines a provably secure operating system kernel which enforces DoD security regulations upon computer computations. The initial implementation of the Security Kernel for KSOS shall execute upon a Digital Equipment Corporation PDP-11/70 computer system. The minimum, typical and maximum hardware configurations to be supported by KSOS are described in the System Specification [A-Specs 78].

3.1.1 Interface Requirements

The KSOS Security Kernel shall have the following interfaces between major system components:

- a. between all software elements and the host computer hardware;
- b. between all software executing in the supervisor and user virtual memory domains of the PDP-11/70;
- c. between the Kernel and all privileged security related subsystems defined in the Non-Kernel Security Related Software CPCI [NKSR 78].
- d. between the user terminal and the Security Kernel.

All information exchanges between user programs, privileged subsystems, files, and devices shall be mediated by the KSOS Security Kernel.

3.1.1.1 Interface Block Diagram

Figure 3-1 is the interface block diagram for the KSOS Security Kernel. This diagram illustrates the layers of virtual machines or functional capability of KSOS. The Security Kernel interfaces directly with the KSOS UNIX Emulator CPCI, user programs, and the host (PDP-11/70) computer. Both the Emulator and the user programs shall communicate directly with the Kernel through an inter-domain call instruction such as the PDP-11 EMT instruction. To preserve compatibility with existing UNIX applications on the PDP-11/70, the PDP-11 TRAP and IOT instructions shall be reserved for the emulation of the UNIX operating system. The Security Kernel shall support secure communication paths from the user terminal to the trusted system software.

Figure 3-1. Interface Diagram

3.1.1.2 Detailed Interface Definition

This section further details the interfaces between the KSOS Security Kernel and the software elements listed in Section 3.1.1, Interface Requirements. The objects supported by the Kernel are discussed, followed by a presentation of the structure of processes executing on the Security Kernel. Finally, these concepts are applied to the detailed interface specification between the KSOS Security Kernel and the other KSOS software elements.

Security Kernel B5 Specification

3.1.1.2.1 Kernel Objects

Five objects shall be supported at the Kernel interface: processes, process segments, files, devices and subtypes. The next section discusses the five Kernel objects and the mechanism used to name the objects and to retain control information concerning the object. This discussion is followed by a presentation of the Kernel primitive operations upon these objects.

All Kernel objects shall be referenced by a common designator called the Secure Entity Identifier (SEID, pronounced "seed"). The SEID shall be separated into two parts, the name space partition and the name within the partition. Only one object type shall be assigned to each partition. However, more than one partition may have the same type. The name space partition shall be used to distinguish the object type: file, device, process segment, file subtype, and process. The name space partition for a file designates the logical volume ("extent") on which the file resides. The System Administrator should make sure that the logical volumes that could be mounted on his system(s) have unique name space partitions. The Immigration Officer function of the Non-Kernel Security Related Software [NKS 78] aids in this task by maintaining a data base of name space partitions in use. The name within the partition shall be used to specify particular instances of an object within the partition. A SEID shall be returned as the result of new object creations (i.e. K_create, K_build_segment, K_fork, and K_spawn.)

Associated with each Kernel object are two classes of information:

- a. type independent information: information common to all objects. The type independent information includes:
 1. the security level of the object (classification category and compartment set)
 2. the integrity level of the object (classification category and compartment set)
 3. the discretionary access control information. As in UNIX, this information shall consist of three sets of permissions:
 - ♦ for processes with the same user ID as the object (in UNIX "the owner"),
 - ♦ for processes with different user IDs, but the same group ID ("the group"), and
 - ♦ for processes with both different user IDs and group IDs ("all others").

Each permission shall consist of three independent access modes:

- ♦ read,
- ♦ write, and

Security Kernel B5 Specification

♦ execute/search.

4. the object's owner (a user ID and a group ID)
 5. the set of privileges associated with the object.
- b. the type dependent information: information whose content varies by the type of the object

The type independent information for an object contains all of the information required by the Kernel to make security decisions (i.e. to permit or deny an access attempt by a process to an object). Because the type independent information is the same for all objects, the security checking may be implemented more simply. As is shown in the formal specifications for the Kernel (Appendix A to this specification), all mandatory access checking can be localized in a single function, SMXflow, which checks whether information may legitimately flow from the requested source to the requested destination.

The discretionary access checking can be similarly localized in a single function, SMXdap. The Kernel shall normally enforce only the read and write permissions. However, some processes possess the privilege to use the execute/search permissions for read permission ("realize execute permissions"). This privilege is used to faithfully emulate UNIX execute-only files (by the Process Bootstrapper CPC) or search-only directories (by the Directory Manager CPC).

3.1.1.2.1.1 Kernel Object: Processes

In KSOS the process is the only active agent in the system. All the other Kernel-supported objects are passive. The type independent information for the process and for the object being accessed shall be used by the Kernel to determine whether to permit or deny access attempts by the process. A process shall have two owner identifications: the effective owner (user ID and group ID) and the real owner. The effective owner shall be used in discretionary access checking, and shall be inherited by new objects created by this process. Thus the effective owner is part of the type independent information. The real owner is part of the type dependent information for the process.

Because a process is itself an object, the Kernel shall enforce the mandatory and discretionary access policies on accesses to the process. In particular, interprocess communication (IPC) shall only be permitted if information may flow from the sending process to the recipient.

The type dependent information for a process shall include:

- a. the identities (SEIDs) of this process and its parent process (the parent process is the process which brought this one into existence via a K_fork or a K_spawn Kernel call)
- b. the process family name
- c. the real owner (user ID and group ID)

Security Kernel B5 Specification

- d. the program counter (pc) (PDP-11 specific)
- e. the program status word (ps) (PDP-11 specific)
- f. the pseudo interrupt level
- g. the advisory priority
- h. the timer alarm value (causes a pseudo interrupt when it counts down to zero)
- i. the supervisor mode execution time (PDP-11 specific)
- j. the user mode execution time
- k. the timer toggle TRUE => pseudo interrupt every clock tick FALSE => no interrupt

In order to provide an operationally viable system, the Kernel shall provide support for privileged processes. Privileged processes shall be either created at system initialization or become privileged through the execution of the Kernel calls `K_invoke` or `K_spawn` specifying that a privileged instruction segment is to replace the segments of the process. The privileges of a process shall persist until another `K_invoke` is issued by the process, at which time the new privileges are established (`K_invoke` can also remove any special privileges for the process) or until a suitably privileged process explicitly alters them through a `K_set_object_level` call. The privileges that may be granted to a process are described in 3.1.1.2.5.

The UNIX Emulator [Emulator 78] (with the exception of the Directory Manager) and all UNIX user domain programs shall have no special privileges. Thus, the Emulator may be modified by KSOS sites to provide locally desired UNIX calls. Naturally, care must be taken in such modifications so that the remainder of the UNIX interface provided by the Emulator is unaffected. However, such changes cannot affect the security properties of the system, because the Emulator is not "trusted" by the Kernel, and cannot violate any of the security rules of the system.

All privileged software shall communicate directly with the Kernel. That is, such software shall use the Kernel calls directly and shall not use the (untrusted) UNIX Emulator. If privileged software is modified, it must be subject to the same rigor (e.g. formal specification, etc.) that were employed in its initial creation. Thus, most KSOS sites will not be allowed to modify or create privileged software.

To the KSOS Kernel, a process shall occupy multiple, separate virtual memory domains. In the PDP-11/70 KSOS implementation, a process shall be composed of three parts:

- a. User domain process segments
- b. Supervisor domain process segments

c. Processor state (registers, etc.)

These Kernel objects shall be visible to programs operating within the supervisor and user domains. The process may also request the Kernel to set or get its type independent and type dependent information via the appropriate Kernel calls.

3.1.1.2.1.2 Kernel Object: Process Segments

A process segment shall be a portion of the virtual memory of a process. A process shall be capable of having a subset of its active segments actually mapped into its virtual address spaces. Process segments shall be variable sized. A single segment shall be capable of spanning a virtual address range from a hardware limited lower bound up to the entire virtual address space in a domain. The development contractor may elect to provide segments in only a pre-defined set of sizes. All process segments, regardless of domain, shall be manipulated by the same set of Kernel segment primitives. User domain segments shall be directly addressable to the Emulator (e.g. through the PDP-11/70 MFPI(D) instructions) in order to transfer operating system call arguments. Programs executing in the user domain shall not directly observe or modify the supervisor domain. User and supervisor programs shall not directly observe or modify Kernel memory. In the general, machine independent case, domains of higher privilege shall be able to observe and to modify domains of lesser privilege. Domains of lesser privilege shall not have any access rights to domains of greater privilege.

Like the other Kernel objects, each segment shall have type independent information associated with it. This information is used by the Kernel to mediate access attempts to the segment. Specifically, the Kernel shall use the type independent information in determining whether other processes may map the segment into their virtual address space and in setting up the hardware memory mapping/protection registers for the processes permitted access to the segment. The privileges contained in the segment's type independent information shall be used by the K_invoke and K_spawn mechanisms. The process privileges shall be set from the privileges of the intermediary specified in the call. The type dependent information for segments shall include the following:

- a. Sharable | Non-Sharable (This information supplements the discretionary access information. A segment might be marked as readable and writable by its owner, but still not be sharable to other processes with the same owner, for example the data segments of a process.)
- b. Swappable | Locked in physical memory.
- c. Status when the reference count drops to 0 (destroy segment|save segment) (the UNIX "sticky bit", ISVTEXT)
- d. Direction of segment growth (upward | downward)
- e. Advisory position for the segment when primary memory resident
- f. Executable | Not Executable (This is more generic. In the PDP-11/70 case, the segments placed in data space may be executable also.)

Security Kernel B5 Specification

- g. the direction of growth for the segment (Up | Down)
- h. the segment size

This information and the type independent information for the segment shall be used by the Kernel in determining whether to permit or deny attempts by processes to access the segment (map it into their address space). For permitted accesses, this information shall be part of the information used to set up hardware protection and/or memory management registers for processes mapping the segment into one of their virtual address domains. For each process mapping the segment into one of its virtual address domains, that use of the segment shall have associated with it:

- a. the domain into which the segment may be mapped
- b. the virtual address in that domain at which the segment starts
- c. whether the segment resides in the Instruction or Data spaces (for the PDP-11/70 implementation.)
- d. whether the segment is presently active (mapped into the virtual address space)

These attributes shall control the use of the physical memory allocated to the segment. This information may be returned with the other type dependent information about the segment.

The concept of KSOS process segments is intended to be machine independent. No assumptions are made about the allocation or use of the native memory management hardware. Only two effects of the underlying hardware shall be visible to the Kernel user:

- a. When a segment is allocated, a new mapping register(s) shall be allocated in a manner such that all segments shall be individually mapped and unmapped from the physical segmentation hardware.
- b. Attempts to map a segment into a process virtual address domain shall fail when the pool of available memory management resources (i.e., registers) or virtual memory space has been exhausted or when there would be overlap of the virtual memory spaces of two segments.

The intent of this approach is to maintain a high degree of machine independence.

In the PDP-11/70 implementation of KSOS as many as sixteen memory management registers shall be available for segment mapping per process domain memory space. If the Instruction and Data (I and D) space separation feature of the PDP-11/70 hardware is not utilized, only eight registers shall be available to a process within a domain. The maximum length of a memory vector supported by a PDP-11/70 map register is 4096 words. In order to create process segments that are longer than 4096 words, more than one memory management register shall be used. (It should be recognized that in order to span the maximum virtual memory of the PDP-11/70, segments must be an integral multiple of 4096 words and must

Security Kernel B5 Specification

be aligned on 4096 word boundaries.) The ultimate size of a KSOS process will be limited by the addressing capability of the PDP-11/70; 64K words if I and D separation is enabled, 32K words if it is disabled. Standard UNIX process segments can be emulated by the following combinations of attributes:

a. Program pure text:

1. Separated (in UNIX, a "411 magic" number): Sharable, swappable, upward growing, instruction space placement, executable, read only
2. Non-separated (in UNIX a "410"): Sharable, swappable, upward growing, data space placement, executable, read only

b. Data, BSS and impure text: Non-Sharable, swappable, upward growing, data space placement, Non-executable, read/write

c. Stack: Non-Sharable, swappable, downward growing, data space placement, Non-Executable, read/write (In addition, UNIX conventionally starts the stack at the highest virtual address and expands it downward.)

d. IPC segments: Sharable, swappable, upward growing, data space placement, Non-Executable, read/write (or read-only)

Other combinations would be possible. For example, it may be possible to implement multiple shared text segments, permitting the construction of the common subroutine mechanism as suggested in the proposals for UNIX IPC [Holmgren et al. 77] [Nemeth et al. 77] [Sunshine 77] [Zucker 77]. Existing UNIX programs which do not use shared IPC segments shall find the KSOS process segment limitations identical to those of standard UNIX.

3.1.1.2.1.3 Kernel Object: Files

The file system visible at the Kernel interface shall be a "flat" file system. No structure shall be imposed upon the organization of files. File directories will be supported by the UNIX Emulator and the Directory Manager. Hence, the file objects at the Kernel level shall be homogeneous. The Kernel files may be thought of as linear sequences of data blocks. The Kernel shall place no special interpretation on the contents or format of the files, e.g. there is no Kernel support for more sophisticated access methods.

The type dependent information associated with a file shall include:

- a. the size of the file
- b. the link count for the file
- c. the time of the last modification to the file
- d. the subtype (if any) to which the file belongs
- e. whether or not the file is open for writing (or was open when the system halted)

Security Kernel B5 Specification

A process shall be able to access the file data contents only if the multilevel security/integrity conditions are satisfied and both the file subtype and file discretionary permissions given to the effective group and user ID of process properly match the requested file operation.

Files shall be referenced at the Kernel level through the use of the Secure Entity Identifier (SEID) that uniquely identifies the file. Access to a file shall be granted by the Kernel before any input or output may be performed to the file. The operation of opening a file shall return an open object descriptor in a manner similar to the current Standard UNIX file open. The open object descriptor shall be used to reference the file in all subsequent Kernel calls. (The operation of opening a file may be viewed as being equivalent to granting a capability to the data object in a capability based Kernel.)

3.1.1.2.1.4 Kernel Object: Devices

As in UNIX, in KSOS devices may be thought of as a special class of files. The type dependent information for devices shall include the same information as the type independent information for files. The type dependent information for devices may include additional information used internally by the Kernel. Devices SEIDs shall be permanently assigned in the Kernel. However, any UNIX directory entry may point to a device. Because the underlying devices upon which the file system is built are visible at the Kernel interface (i.e. they may be opened by suitably authorized processes), the System Administrator should assure that these devices are protected from unauthorized modification or access by means of appropriate discretionary access and subtype membership. The Kernel shall not allow files to be created whose security level is greater than the underlying device on which they are being created.

3.1.1.2.1.5 Kernel Object: File Subtypes

File subtypes are a type extension feature. The goals of the KSOS Security Kernel file subtype mechanism are to:

- a. Subdivide the Kernel file object type
- b. Provide a separate subtype manager for each file subtype controlling the use of specially typed files (e.g. directories)
- c. Localize subtype management improving leverage on the proof process

File subtypes, because they are Kernel objects, shall have both a Secure Entity Identifier (SEID) and a type independent information. Thus, file subtypes shall have a security and integrity level associated with each instance of a subtype. Most importantly, file subtypes shall have an owner and a discretionary permission list. The subtype owner shall refer to the group and user ID of the subtype manager. The discretionary permission list shall describe the permitted access modes to all files of that subtype by the subtype manager (owner), group, and all users.

The nature and structure of file subtypes is best revealed by examining their use in the Kernel level file open function, `K_open`. In order to open a file for access in a specific mode, the process shall successfully meet four

Security Kernel B5 Specification

classes of conditions:

- a. Mandatory security and integrity
- b. Subtype capability or file subtype discretionary permission
- c. File data discretionary permission

The subtype capability of a process may be either implicit or explicit. Implicit subtype capability shall be determined by the effective group and user ID of the process. If the requested access mode successfully matches the permissions accorded to the effective group and user ID of the process for that file subtype, the file shall be opened. Explicit subtype capability shall be acquired by performing a `K_open` call on the subtype, specifying the subtype by its SEID. Subsequent `K_open` calls applied to files belonging to this subtype shall accept the open object descriptor returned by the subtype open operation as an explicit capability (permission) to manipulate the file in the requested mode.

As an example of the use of file subtypes, the Emulator uses the Kernel subtype mechanism to implement UNIX directories. Providing that they meet the mandatory and discretionary access checks, ordinary users are permitted read access to directory files, but are denied write access. Only the Directory Manager is permitted to have write access to directory subtype files. Hence, all user processes are able to search directories and perform their own path name interpretation. However, only the Directory Manager is able to update or delete directories, insuring the integrity and accuracy of directory files.

3.1.1.2.2 Process Structure

The logical information which determines a specific instance of a process, in the PDP-11/70 implementation of KSOS, shall span three memory domains. The data contained at each domain level is characterized by the function it is intended to serve:

- a. User domain: User process information (i.e., program instructions, data, stack, etc.)
- b. Supervisor domain: Information required to faithfully emulate the UNIX operating system. Further, all Non-Kernel Security-Related Software functions which operate with special privileges or which manipulate security critical data bases shall utilize only Kernel calls (i.e. these functions shall not utilize the Emulator) and shall reside in the Supervisor domain.
- c. Kernel: Process data needed to multiplex the hardware, maintain the logical determinacy of user processes, and to enforce multilevel security.

Unprivileged processes executing in either user or supervisor domain shall be regarded as "untrusted" by the Kernel domain software. That is, the security, integrity, and discretionary access policies shall be enforced by the Kernel for all unprivileged processes.

Security Kernel B5 Specification

3.1.1.2.3 Hardware Interface

The KSOS Security Kernel shall execute correctly on all hardware configurations specified in the System Specification (Type A) [A-Specs 78].

The following general remarks can be stated about the KSOS hardware requirements. The KSOS design shall execute on a host computer with a minimum of three separate virtual memory and privilege domains. The KSOS Security Kernel shall reside in the most privileged host processor domains. Non-Kernel software shall reside in processor domains where all accesses to Kernel-supported system resources can be mediated. This means that non-Kernel software may not change its virtual to physical address mapping, potentially gaining access to information resident in other processes, files, or device registers. The processor shall also provide an interdomain call mechanism which permits services to be performed on its behalf by software executing in another virtual memory and privilege domain. That is, the processor shall support some sort of "trap" mechanism for requesting services from more privileged software.

3.1.1.2.4 Emulator/User Mode Software Interface

In addition to the mediation of information accesses, the purpose of the KSOS Security Kernel is to provide a service environment to support applications software executing in Kernel-controlled processes. This specific implementation effort requires the emulation of the Western Electric Corp. UNIX operating system interface as defined in the System Specification. All non-Kernel software shall have a direct call interface to the Security Kernel. Hence, any non-Kernel software shall be able to directly utilize Kernel level services. Briefly, these Kernel level services are the manipulation of system resources, i.e. files, devices, processes, process segments, and file subtypes. Section 3.2 of this document details the Security Kernel functions available through the direct Kernel call interface. Applications programs may include both UNIX calls and Kernel calls. However, the application program designer must take care that the use of both types of calls does not lead to incorrect results due to interactions between the Emulator's use of Kernel calls, and that of the application program. In no cases can such incorrect interaction lead to security, integrity, or discretionary access violations. Rather, the potential problem is in incorrect results or garbling of files.

In order to accurately emulate the UNIX operating system interface on the PDP-11/70 host processor, the TRAP instruction shall be used to transfer user domain requests to the UNIX Emulator operating in the supervisor domain of the machine. The IOT instruction shall also cause a trap to the Emulator for the invocation of core dumps by the Emulator at the request of the user domain program. The EMT instruction shall be reserved for direct user domain to Kernel and supervisor domain to Kernel interdomain calls. Parameters to Security Kernel calls should (in the PDP-11 implementation) be passed on the memory stack of the calling domain. To facilitate the transfer of lengthy blocks of data (i.e. K_device_function, K_read_block, and K_write_block) an address pointer (called a "parameter block") to the data area may be used as a valid parameter. The Kernel should use this address to fetch/store information from the data area. Return values from the Security Kernel should be transferred to the caller in a register (i.e. register R0) where possible. When the Kernel is requested to transfer data into or out of the user's memory, the data to be transferred must be

Security Kernel B5 Specification

completely contained within one process segment. This shall include the parameters passed in or returned on the stack.

The KSOS Security Kernel shall reflect certain processor traps to the Supervisor domain process, i.e., the UNIX Emulator. The processor exception conditions reflected outwards shall be those traps caused by the software executing outside the domain of the Kernel. The PDP-11/70 traps to be reflected outward by the Kernel are:

- a. Illegal or reserved instruction exceptions.
- b. Breakpoint traps.
- c. IOT instruction traps. (Causes UNIX abort.)
- d. TRAP instruction. (Causes entry to UNIX emulator dispatch.)
- e. Floating point errors.
- f. Memory management traps.

The means through which these traps shall be communicated to the executing process shall be a pseudo interrupt. The pseudo interrupt may include additional information, such as the cause of the trap, or any relevant processor status required by the extra-Kernel software to process the trap.

3.1.1.2.5 Privileged Subsystem Interfaces

To provide a useful, operating environment, some non-Kernel software must be privileged (trusted) to perform system maintenance and other functions that selectively violate the protection policy of the Kernel. The KSOS Security Kernel shall maintain the process privilege information which determines the policies that the process has been permitted to violate. Privileged software shall be brought into execution either at system initialization time or through the execution of the `K_invoke` or `K_spawn` Security Kernel primitives. An orderly system initialization procedure shall be implemented which initializes the trusted processes required for correct operation of the overall system. The `K_invoke` (or `K_spawn`) primitive shall derive the process privileges from the designated instruction segment in a manner which guarantees the execution of that segment with the privilege set. One particular instruction segment, containing the Process Bootstrapper CPC, shall have the privilege to set privileges, and shall be a mechanism by which a process begins executing a potentially privileged image contained in a file. The Process Bootstrapper shall set the privileges of the process to those of the file from which it was initialized. The privilege information associated with the process image file is normally set only by the Privilege Control Editor which may only be executed by the System Administrator or System Security Officer while running at the ADMINISTRATOR integrity level. (See the Non-Kernel Security Related Software B5 Specification [NKSr 78].)

The following special privileges shall be capable of being granted or revoked:

Security Kernel B5 Specification

- a. The ability to violate the simple security property. (This privilege is not presently used. It is included for completeness.)
- b. The ability to violate the simple integrity property.
- c. The ability to violate the *-property for security.
- d. The ability to violate the *-property for integrity.
- e. The ability to violate the security compartment checking rules.
- f. The ability to violate discretionary access. (This privilege is not used. It is included for completeness.)
- g. The ability to use the execute discretionary access permissions for read access. (This is also called "realize execute permissions".)
- h. The ability to use the `K_secure_terminal_lock` primitive.
- i. The ability to use the `K_link` and `K_unlink` primitives.
- j. The ability to use the `K_mount` and `K_unmount` primitives.
- k. The ability to change the following type independent information about an object:
 - 1. The security level of the object.
 - 2. The integrity level of the object.
 - 3. The owner (user ID and group ID) of the object.
- l. The ability to change the privileges in the type independent information for the object.
- m. The ability to set type dependent information (status) about an file (i.e., the ability to execute `K_set_file_status`).
- n. The ability to save a segment in the swapping area (i.e. the ISVTEXT or "sticky bit").
- o. The ability to lock a segment in core (not swappable).
- p. The ability to use the `K_signal` primitive.
- q. The ability to use the `K_walk_process_table` primitive.
- r. The ability to use the `K_halt` primitive.
- s. The ability to use the Kernel primitives from user mode. (This privilege is included for protection of the Emulator. It shall be possible to disable it, allowing unrestricted use of Kernel calls from user domain.)

It should be noted that the owner of an object may always change the discretionary access information for that object, hence no special privilege is needed.

There are some privileges which subsume others. For example, the ability to set the level of an object (change its type independent information) is functionally equivalent to being able to violate the simple security, security *-property, simple integrity, integrity *-property, and discretionary access rules. In general privileged processes should be afforded the minimum privileges that allow them to perform their intended functions. Programs which need the more powerful privileges, e.g. set object level, or set privileges, should be subject to the most design and implementation rigor.

3.1.1.2.6 Secure Terminal Interface

Certain privileged processes such as the Secure Server (part of the Non-Kernel Security-Related Software CPCI [NKSr 78]) require a secure communication path from the user terminal to transfer passwords and other information which cannot be contaminated or compromised. Further, the user must have the assurance that he is communicating solely with the Kernel or trusted subsystem and not a user process masquerading as a trusted program. The Secure Terminal Interface function of the KSOS Kernel shall provide this mechanism.

Two sets of terminal device SEID's shall be maintained, the user (thawed) set and the secure (frozen) set. The user set shall determine the access capabilities of regular user processes to specific terminals. The secure set shall determine the access capabilities of secure software to specific terminals. The secure set shall be protected by the discretionary access mechanism. Their owner shall be a distinguished user ID, and software allowed to use them shall execute in the discretionary access domain of this user ID. The terminal is said to be "frozen" when the secure path is active and the normal path is blocked. It is said to be "thawed" when the normal path is active. A terminal shall become frozen when the user enters the secure attention character, a character reserved for this one purpose. When the Kernel receives the secure attention character it shall freeze the terminal and send an IPC message to the Secure Initiator [NKSr 78]. A terminal shall become thawed only by having a privileged process issue a K_secure_terminal_lock Kernel call. When a terminal is frozen, all input and output on the normal path shall be blocked. It is the responsibility of the processes privileged to thaw terminals to assure that only the correct path is thawed, i.e. that the processes using that path have the appropriate security and integrity level to use the terminal at its (possibly new) level.

In the frozen mode, only software privileged to use the secure path may interact with the user. (The exact privileges of such software shall be the privilege to use K_secure_terminal_lock and the ability to execute in the discretionary access domain user ID which owns the secure terminal devices.) In order to communicate with a frozen terminal, a privileged process shall open the terminal using the corresponding terminal SEID selected from the secure terminal set. Only privileged software shall successfully open terminals using SEID's from the secure terminal set. The returned open object descriptor may then be used by the secure process to perform read and write operations to and from the terminal.

3.2 Detailed Functional Requirements

The functional descriptions in this section use the same argument names and order as the formal specifications found in Appendix A. However, these descriptions have expanded out various structures into their components for clarity and flexibility of implementation. Subject to Government approval, the development contractor may package the arguments as needed for ease of implementation and use.

Many of the following functions have a list of exception conditions which must be satisfied for the function to have any effect. The order of these exceptions in the text is not necessarily the order in which they should be checked in an implementation of the Kernel. Rather, the order of the exceptions in the formal specifications found in Appendix A shall be used. The development contractor may add additional exceptions, subject to Government approval. In some cases, the true cause of an exception cannot be returned to the user because it could be used as a confinement channel [Lampson 73] [Lipner 75]. In these cases, both the true cause and the returned value are noted.

Some of the functions utilize parameter blocks to designate areas in which to accept or return data. A parameter block is a structured argument that specifies an area of virtual memory for the process. For the PDP-11/70 implementation, a parameter block shall consist of:

- a. the domain in which the data resides (user or supervisor)
- b. the space in which the data resides (instruction or data space)
- c. the starting virtual address of the data
- d. the size of the data block (this may be omitted if a standard size is always used)

In all cases, an exception shall be raised if the data area specified by the parameter block does not lie completely within one segment which is accessible to the caller. That is, the segment containing the data specified by the parameter block must be capable of being accessed in the implied mode. For data that is being fetched, the segment must be readable, and for data being returned, the segment must be writable. Where no ambiguity results, the phrase "contained in the parameter block" (or similar) is used in lieu of the more precise "contained in the data area designated by the parameter block" in the interests of brevity.

3.2.1 Kernel Function: K build segment

3.2.1.1 Inputs

The Kernel primitive K_build_segment shall take the following parameters.

- ss.gl.shareable: Boolean flag for sharing FALSE => segment is not sharable, TRUE => sharable
- ss.gl.lock: Boolean flag for locking TRUE => do not swap (segment is locked in memory) (requires privilege to lock a segment) FALSE => swappable
- ss.gl.sticky: Boolean flag for segment status when reference count goes to

Security Kernel B5 Specification

zero (the "sticky bit" of UNIX) TRUE => do not destroy segment
(requires privilege to set the sticky bit) FALSE => destroy segment
(normal case)

ss.gl.memAdvise: Boolean flag designating advisory information to the Kernel
regarding placement of the segment in memory, TRUE => segment will have
I/O directed to it, FALSE => no I/O. This is merely advisory, I/O can
be done to any segment accessible in the required mode.

ss.gl.executable: Boolean flag indicating whether this segment is executable

ss.gl.growth: the direction of segment growth (up or down)

ss.size: the initial segment size in bytes (the development contractor may
round this up to the next integral unit of allocation)

ss.access: access allow calling process to the segment

ss.vl.domain: domain of allocation/mapping

ss.vl.idSpace: Instruction Space or Data Space placement (PDP-11/70
specific)

ss.vl.vAddr: virtual address of first byte in the segment

ms: the initial discretionary access permissions (read, write,
execute/search for the owner, group, and all others)

3.2.1.2 Processing

K_build_segment shall create a new process segment and map it into the process. Memory management registers shall be used by the segment according to the address range specified for the segment. An error shall be returned if the address range specified for the segment would be in conflict with the address range of some other mapped in segment.

The discretionary access parameter shall determine the type of machine reference permitted to the segment and which other processes may share the segment if sharing is permitted. Accesses shall be selected from the set {read, write, execute}. The segment shall be initialized to contain all zeros. The process must store whatever data it requires into the segment explicitly. The owner of the segment shall be taken from the effective owner of the process. The security and integrity levels of the segment shall be taken from those of the process.

Segments may be shared between processes providing that the security, integrity and discretionary access checks would allow such sharing. Sharing of segments requires that the first process to desire to share the segment create it. Subsequent to creation, the segment SEID must be made available to processes wishing to share the segment, typically either via placing it in a mutually agreed upon file, or by passing it via an IPC message. The other processes would then issue K_rendezvous_segment calls (below) to gain access to the segment.

It is the responsibility of the processes sharing the segment to see to it that it is properly initialized, either by the Kernel's guarantee of all zeros, or by explicit initialization. The initialization may require cooperation or mutual exclusion to be completed successfully. This is particularly true in the case of shared pure text segments which are to be resident in the Instruction Space on the PDP-11/70. Since such segments cannot be written into, they must be created as writable segments, initialized from the appropriate image file,

Security Kernel B5 Specification

and then have their status changed to make them reside in the correct space and be non-writable. Care should be taken in the design of programs like the Process Bootstrapper which perform such initializations to assure that duplicate initialization attempts or multiple copies of the same shared text segment do not occur.

The virtual address and domain parameters shall be for the calling process only. Other processes sharing the segment may map it into their virtual address spaces as desired through their use of the `K_rendezvous_segment` calls. Of course, some segments, e.g. text segments, may operate correctly only if mapped starting at a specific virtual address.

3.2.1.3 Outputs

The Kernel primitive `K_build_segment` shall return the following information.

ec: error conditions
result.segSeid: the seid of the created segment
result.segd: designator name (process local) of newly created segment

The Kernel primitive `K_build_segment` shall detect and return the following exception (error) conditions.

EX: insufficient privilege, attempted to make new segment either sticky or locked in memory without the appropriate privilege
EX: bad discretionary access mode, attempted to create a segment which was writable but not readable for some permission (owner, group or other)
EX: bad size, segment size was too large
EX: global resource exhaustion (this is a potential channel)
EX: too many active segments
EX: virtual memory error, virtual memory address space would be exceeded or overlapped, or would fail to touch an 8K byte boundary

3.2.2 Kernel Function: K close

3.2.2.1 Inputs

The Kernel primitive `K_close` shall take the following parameters.

od: open object descriptor

3.2.2.2 Processing

The file referenced by the `K_close` call shall be inaccessible until the file has been successfully opened again. If the object associated with the open descriptor, od, is a file whose link count has been decremented to zero, and there are no other opens to it (i.e. no other open descriptors in the system are associated with it), then the file shall be deleted as part of the `K_close` processing.

Security Kernel B5 Specification

3.2.2.3 Outputs

The Kernel primitive `K_close` shall return the following information.

`ec`: error condition

The Kernel primitive `K_close` shall detect and return the following exception (error) conditions.

EX: open object descriptor does not specify an open object

3.2.3 Kernel Function: K_create

3.2.3.1 Inputs

The Kernel primitive `K_create` shall take the following parameters.

`nspSeid`: SEID from the name space in which to create file

`ms`: discretionary access permissions

`stCap`: subtype capability, an open descriptor for a subtype. The newly created file will belong to that subtype if the call is successful.

3.2.3.2 Processing

`K_create` shall be the mechanism by which new files are created. `K_create` differs from the UNIX `creat(II)` call in that `K_create` always creates a new file. `K_create` shall attempt to create the new file in the name space partition of the SEID supplied as an argument. The file control information shall be allocated and initialized. The Kernel level file name (SEID) shall be returned for potential inclusion in an Emulator level directory or similar construct. In addition, an open object descriptor shall be returned. This shall be equivalent to the user having requested a `K_open` for writing on the newly created file, with the exception that the discretionary access permissions are not checked. (This is to allow the common UNIX construct of lock files created with no access permissions.) The discretionary access shall be taken from the supplied argument. The owner of the file shall be taken from the effective owner of the process. The security and integrity levels of the file shall be taken from those of the process. The newly created file shall belong to the subtype supplied via the subtype open descriptor, `stCap`, providing that `stCap` represents a valid subtype open for writing.

3.2.3.3 Outputs

The Kernel primitive `K_create` shall return the following information.

`ec`: error conditions

`result.fSeid`: secure entity identifier

`result.od`: open object descriptor

The Kernel primitive `K_create` shall detect and return the following exception (error) conditions.

Security Kernel B5 Specification

EX: security/integrity level of file system do not permit object at this level to be created on it
EX: too many open files
EX: the prototype SEID supplied is not for a file
EX: file system is read only
EX: file system not mounted
EX: subtype invalid, the subtype open descriptor supplied was not for a subtype open for writing
EX: unable to create file (this is a possible confinement channel)

3.2.4 Kernel Function: K device function

3.2.4.1 Inputs

The Kernel primitive K_device_function shall take the following parameters.

od: open object descriptor
f: function requested
arguments.vloc.domain: domain of parameter block containing input arguments
arguments.vloc.idSpace: space of argument parameter block
arguments.vloc.vAddr: address of argument parameter block
arguments.size: size of argument parameter block (may be omitted and have parameter block default to a standard size)
status.vloc.domain: domain of parameter block in which status will be returned
status.vloc.idSpace: space of status parameter block
status.vloc.vAddr: address of status parameter block
status.size: size of status parameter block (may be omitted and have parameter block default to a standard size)
id: asynchronous call identifier

3.2.4.2 Processing

The K_device_function primitive shall permit the process to perform special device control operations. The device shall be opened prior to the K_device_function call. The device shall be referenced only through an open object descriptor. Hence, authorized access to the device shall always be guaranteed. In addition, the requesting process must own the device. That is, the effective owner of the process must be the same as that of the device, or the process must have the privilege to change its owner. The arguments parameter block passed to the K_device_function call shall designate an area of the user process memory which contains parameters required to perform the specific function, f, requested. Thus, the particular parameters may vary depending upon which function was requested. The device specific routine shall be responsible for verifying the validity of argument parameters. In particular, certain functions may require that the process be privileged. The status parameter block shall designate an area of the process virtual memory to be used to return device specific status information. Thus, the particular parameters may vary depending upon which function was requested.

Security Kernel B5 Specification

The access checking on the two parameter blocks shall be similar to that performed by the `K_read_block` call (below) for the arguments parameter block, and to that performed by `K_write_block` (below) for the status parameter block. Error returns from this call shall include invalid file descriptor and handler dependent parameter value errors in addition to the checks on the validity of the virtual addresses of the parameter blocks, the validity of the open object descriptor, and that the user owns the affected device.

Limited numbers of asynchronous `K_device_function` requests may be made. If the asynchronous id field is non-null, `K_device_function` may return before completion of the requested operation. At the completion of the requested operation, the Kernel shall transmit a device completion pseudo interrupt, similar to that transmitted at the completion of `K_read_block` or `K_write_block`. The values returned in the area designated by the status parameter block may not be valid until the pseudo interrupt has been transmitted.

3.2.4.3 Outputs

The Kernel primitive `K_device_function` shall return the following information.

ec: error conditions

status: status data in status parameter block

The Kernel primitive `K_device_function` shall detect and return the following exception (error) conditions.

EX: invalid open object descriptor

EX: descriptor does not refer to a device

EX: user does not own device

EX: invalid parameter block (either arguments or status)

EX: invalid device operation (command)

EX: invalid device dependent parameters in area designated by arguments parameter block

3.2.5 Kernel Function: K fork

3.2.5.1 Inputs

The Kernel primitive `K_fork` shall take the following parameters.

None

3.2.5.2 Processing

The `K_fork` primitive shall create a process. The new process shall be remembered as a child of the caller (parent). The child shall be an exact duplicate of the parent. The process id of the parent shall be returned to the child. The process id of the child shall be returned to the parent. The program counter of the parent is NOT adjusted as it is in standard UNIX. The returned process id will be sufficient to distinguish parent and child. The

Security Kernel B5 Specification

type independent information of the child process is identical to the type independent information of the parent process.

The non-sharable segments of the process shall be copied into new segment instances for the child. The reference counts of sharable segments shall be incremented. The process local segment names shall be the same in both the parent and child, although in the case of non-sharable segments they shall refer to a different segment instance (and therefore, a different segment SEID) for the child than for the parent.

The child shall inherit the open objects of the parent. That is, each object that is open in the parent shall be opened in the child, and shall have the same local open object descriptor. The open counts of the open objects so inherited shall be incremented to reflect the fact that another process (the child) has them open. If the parent has an object open for exclusive use, the K_fork call shall fail, preventing two processes from having simultaneous access to the same exclusive use object.

An error shall be returned if the number of available processes has been exhausted. Analysis will be required to determine the bandwidth of this binary information channel and to select an appropriate solution.

3.2.5.3 Outputs

The Kernel primitive K_fork shall return the following information.

childSeid: process SEID (the parent gets childSeid, and the child gets the process SEID of the parent)
ownSeid: process SEID (the parent gets parentSeid, and the child gets childSeid)
ec: error conditions

The Kernel primitive K_fork shall detect and return the following exception (error) conditions.

EX: exclusive use files open
EX: unable to create new process (possible information channel)

3.2.6 Kernel Function: K_get_file_status

3.2.6.1 Inputs

The Kernel primitive K_get_file_status shall take the following parameters.

fSeid: file seid

3.2.6.2 Processing

K_get_file_status shall return the type dependent information associated with a file. The exact format and information content of the file status block shall be left to the discretion of the development contractor subject to

Security Kernel B5 Specification

Government review and approval. The invoking process must have read access to the file with respect to the mandatory security and integrity rules only. No discretionary access checking shall be performed. (This is required for faithful emulation of the UNIX construct of lock files which are created with no discretionary access permissions.)

3.2.6.3 Outputs

The Kernel primitive `K_get_file_status` shall return the following information.

- `ec`: error conditions
- `result.nBlocks`: the size of the file in blocks
- `result.linkCount`: the number of references to the file in higher level constructs (e.g. directories or equivalent)
- `result.timeLastMod`: the time of last modification to the file
- `result.subtype`: the subtype to which the file belongs (a SEID)
- `result.openForWriting`: a Boolean flag indicating that this file is open for writing (TRUE) or was open for writing when the system was halted. Such information is needed for higher order file recovery after unplanned system halts. A value of FALSE shall mean that the file had been normally closed at the time of system shutdown.

The Kernel primitive `K_get_file_status` shall detect and return the following exception (error) conditions.

- EX: object does not exist or is not a file, device or subtype
- EX: attempt to read a higher security level object (the actual return value is as if the object did not exist)
- EX: attempt to read a lower integrity level object (the actual return value is as if the object did not exist)

3.2.7 Kernel Function: K_get object level

3.2.7.1 Inputs

The Kernel primitive `K_get_object_level` shall take the following parameters.

- `seid`: the seid of the desired object

3.2.7.2 Processing

The primitive `K_get_object_level` shall return the type independent information for an object.

3.2.7.3 Outputs

The Kernel primitive `K_get_object_level` shall return the following information.

- `ec`: error conditions

Security Kernel B5 Specification

otii.nd.securityCategory: security classification category of object (e.g. TOP SECRET, UNCLASSIFIED)
otii.nd.securityComp: security compartments of the object
otii.nd.integrityCategory: integrity classification category of the object
otii.nd.integrityComp: integrity compartments of the object (now, always null)
otii.da: discretionary access permissions for the object (read, write, execute/search for the owner, group and others)
otii.owner: object owner (a user ID)
otii.group: object group (a group ID)
otii.priv: the privileges associated with the object

The Kernel primitive `K_get_object_level` shall detect and return the following exception (error) conditions.

EX: object does not exist
EX: attempt to read a higher security level object (the actual return value is as if the object did not exist)
EX: attempt to read a lower integrity level object (the actual return value is as if the object did not exist)

3.2.8 Kernel Function: K_get_process_status

3.2.8.1 Inputs

The Kernel primitive `K_get_process_status` shall take the following parameters.

objSeid: SEID of process about which to get status

3.2.8.2 Processing

The `K_get_process_status` call shall return the type dependent information about a process. A process may successfully request information of processes that it can access given its level and the level of the target process.

3.2.8.3 Outputs

The Kernel primitive `K_get_process_status` shall return the following information.

ec: error conditions
ps.self: the SEID of the process (should be the same as objSeid)
ps.parent: the SEID of the parent process (The parent of a process, p, is the process which issued the Kernel call that brought process p into existence.)
ps.family: the process family identifier
ps.realUser: the real user ID
ps.realGroup: the real group ID
ps.pc: program counter (PDP-11 specific)
ps.ps: processor status word (PDP-11 specific)

Security Kernel B5 Specification

ps.pil: pseudo-interrupt level
ps.advPrio: advisory priority
ps.timerAlarm: the time remaining before a timer interrupt will occur
ps.supervisorTiming: execution time in the supervisor domain (PDP-11 specific) (This need not be an exact timing.)
ps.userTiming: execution time in the user domain (PDP-11 specific)
ps.timTog: Boolean indicating whether repetitive timer interrupts are to be generated

The Kernel primitive `K_get_process_status` shall detect and return the following exception (error) conditions.

EX: process does not exist or is inaccessible

3.2.9 Kernel Function: K_get_segment_status

3.2.9.1 Inputs

The Kernel primitive `K_get_segment_status` shall take the following parameters.

segSeid: seid of the desired segment

3.2.9.2 Processing

`K_get_segment_status` shall return the type dependent information associated with a segment. A process may successfully obtain type dependent status about any segment from which information could flow to the process. No discretionary access checking shall be performed.

3.2.9.3 Outputs

The Kernel primitive `K_get_segment_status` shall return the following information.

ec: error conditions
ss.gl.sharable: Boolean indicating whether or not the segment may be shared
ss.gl.lock: Boolean indicating whether or not the segment is to be locked in memory
ss.gl.sticky: Boolean indicating whether or not the segment is to be retained in the swap area after there are no more active references to it
ss.gl.memAdvise: Boolean indicating whether or not it is expected that I/O will be done to this segment (this is merely advisory, I/O can be done to any segment which is accessible in the desired mode)
ss.gl.executable: Boolean indicating whether or not the segment contains executable data
ss.gl.direction: growth direction for the segment (Up or down) (PDP-11 specific)
ss.size: the size of the segment, in the PDP-11 this is in bytes

Security Kernel B5 Specification

The Kernel primitive `K_get_segment_status` shall detect and return the following exception (error) conditions.

EX: segment doesn't exist or is inaccessible

3.2.10 Kernel Function: K interrupt return

3.2.10.1 Inputs

The Kernel primitive `K_interrupt_return` shall take the following parameters.

none

3.2.10.2 Processing

`K_interrupt_return` shall provide an atomic return operation from pseudo interrupts. It can be thought of as being analogous to the PDP-11 RTI and RTT instructions. When a pseudo-interrupt occurs, the program counter, processor status word, and current pseudo-interrupt level shall be saved in a pseudo-interrupt vector for the particular type of pseudo-interrupt which occurred. In the PDP-11 implementation, these vectors shall be located at fixed locations in the supervisor domain. The `K_interrupt_return` call shall restore the process state from these saved values. Because the interrupted process state is accessible to the process, the `K_interrupt_return` call shall check the saved state prior to restoring it. The process shall not be permitted to increase its privileges or accessible domains. (Similar checking takes place in the processing of the pseudo interrupt itself.) For example, in the PDP-11 implementation, the process shall not be allowed to set the either mode (current or previous) to kernel or to alter the processor priority and general register set information.

3.2.10.3 Outputs

The Kernel primitive `K_interrupt_return` shall return the following information.

none

3.2.11 Kernel Function: K invoke

3.2.11.1 Inputs

The Kernel primitive `K_invoke` shall take the following parameters.

immSeid: segment seid of intermediary process segment
arg: segment designator specifying argument segment

3.2.11.2 Processing

The purpose of the K_invoke primitive is the invocation of potentially privileged software. The effect of this call shall be to replace the existing segment map (including the executing text segment) by a new process image. The segments currently active may or may not be released, at the option of the development contractor, subject to Government approval. The new process image shall have only the intermediary segment and the argument segment active (mapped in). Arguments for use by the intermediary process may be placed in the argument segment. The exact format of the argument segment is determined by the particular intermediary specified in the call. The argument segment may be used by the intermediary as a scratch pad as the intermediary builds any other segments it requires. It shall be the responsibility of the newly executing program (the intermediary) to create its own working segments. The privileges of the process shall be set to those associated with the intermediary segment. In the PDP-11/70 implementation, the intermediary shall be mapped into location 0 of the supervisor domain instruction space, and the argument segment shall be mapped into location 0 of the supervisor mode data space. Thus, the intermediary must be coded as a separated instruction/data space program. The intermediary may perform any arbitrary function. Thus, applications of the KSOS Kernel may elect to create specialized intermediaries to perform specific functions. The only pre-defined intermediary is the Process Bootstrapper, described next.

3.2.11.3 Process Bootstrapper Processing

The Process Bootstrapper segments shall implement a trusted process whose sole purpose is the creation of other, potentially trusted, environments by replacing itself with image from the prototype file whose name is passed in as an argument. The Process Bootstrapper shall have the following privileges:

- a. to set privileges
- b. to set the effective owner
- c. to set the sticky bit
- d. to set the lock (no swapping) bit
- e. to set the security and integrity level
- f. to realize execute permissions (i.e. use the execute permissions for read access attempts)

Using the parameters specified in the argument segment, the Process Bootstrapper shall build a new set of process segments conforming to the process prototype file. The privileges for the new environment shall be obtained from the process prototype file. If the prototype file has no privileges associated with it, the new environment shall be unprivileged. If the prototype file specifies that it is to execute in a different discretionary access domain, the bootstrap shall change the effective user and/or group of the process to the owner of the prototype file. The new trusted process shall then be set into execution by the Process Bootstrapper. Note that a completely trusted path exists from the K_invoke call, through process construction, to the execution of the trusted software.

Security Kernel B5 Specification

The use of the `K_invoke` call shall not be limited to the invocation of trusted processes. Untrusted processes may also be executed through the `K_invoke` call. If change of discretionary access domain or privilege is not required by the type dependent information associated with the process prototype file, the process bootstrap shall simply remove all privileges prior to setting the new image into execution.

The privilege information associated with a process prototype file shall be controlled by the Privilege Control Process, a restricted use program discussed in the Non-Kernel Security-Related Software CPCI Specification [NKSr 78].

3.2.11.4 Outputs

The Kernel primitive `K_invoke` shall return the following information.

ec: error conditions

The Kernel primitive `K_invoke` shall detect and return the following exception (error) conditions.

- EX: invalid argument segment designator
- EX: argument segment is sharable
- EX: argument segment is not writable
- EX: intermediary segment does not exist or is inaccessible
- EX: intermediary segment does not have execute permissions for this user
- EX: argument segment could not fit into memory available to it
- EX: intermediary segment could not fit into memory available to it

The intermediary may detect additional error conditions associated with the particular arguments supplied. Due to the fact that the process environment is discarded prior to the activation of the intermediary, these errors would normally be reported to the parent of the process via an IPC message from the intermediary. The process bootstrap shall report the following exception (error) conditions:

- EX: prototype file is inaccessible
- EX: format of argument segment is invalid
- EX: arguments to invoked process too large

3.2.12 Kernel Function: K_link

3.2.12.1 Inputs

The Kernel primitive `K_link` shall take the following parameters.

fSeid: SEID of file affected

3.2.12.2 Processing

`K_link` shall increment the reference count of a Kernel file. The reference count may be used to indicate the number of UNIX directory entries which point to this SEID. Applications of the KSOS Kernel which do not use the UNIX

directory structure and semantics may use the reference count for other purposes. The reference count may only be incremented by processes privileged to issue this call. Such processes should be carefully designed to reduce the bandwidth of the resultant confinement channels and to preserve the integrity of the higher level directory structure, if any. The security and integrity checking on K_link shall be as if the user were reading and writing the file. No discretionary access checking shall be performed. Thus, the processes privileged to use K_link may implement whatever discretionary checking they chose to.

For example, the Directory Manager of the UNIX Emulator [Emulator 78], when emulating the UNIX link(II) call, checks that the directory in which the new entry is being place is writable by this user prior to making the K_link call.

3.2.12.3 Outputs

The Kernel primitive K_link shall return the following information.

ec: error conditions

The Kernel primitive K_link shall detect and return the following exception (error) conditions.

EX: insufficient privilege

EX: object does not exist or is inaccessible from security or integrity considerations

EX: object is not a file

EX: object is in a file system which has been mounted as read only

3.2.13 Kernel Function: K mount

3.2.13.1 Inputs

The Kernel primitive K_mount shall take the following parameters.

dev: device SEID

leaf: file system new SEID

root: file system old SEID

readOnly: Boolean read only flag TRUE => read only, no writing is permitted on the name space partition being mounted, FALSE => writing allowed

3.2.13.2 Processing

The K_mount call performs two functions:

- a. associating a particular name space partition with a device
- b. causing all references to the "leaf SEID" to refer to a SEID ("root SEID") on the newly mounted name space partition

In KSOS, physical disk volumes may be logically subdivided into "extents". Each

Security Kernel B5 Specification

extent containing a KSOS file system shall belong to a different name space partition. Thus, the SEID of a file is unique across all the file systems that could be mounted. It shall be the responsibility of the System Administrator to assure that the name space partition assignment are unique. The Immigration Officer software [NKSR 78] maintains a data base of name space partitions presently assigned. When an extent is mounted, the Kernel shall update an internal data base that tells it on which device the SEID's in the mounted name space partition may be found. The second facet of the mount mechanism is the logical connection of the newly mounted file system into the rest of the KSOS file system. This shall be accomplished by mapping all references to "leaf SEID", a SEID of an existing, currently available (i.e. the old SEID's file system must itself be presently mounted and the leaf SEID must exist and be inactive) file to a SEID on the newly mounted file system, the "root SEID". It shall be possible to mount a file system as read only, in which case no writing may be done to any file in that file system.

Each extent contains type independent information. The security and integrity levels of the extent shall be interpreted as the maximum levels allowed for any file on the extent. The discretionary access information shall determine who may mount that particular extent. To mount an extent the user must have write permission to it. This shall prevent users from mounting the private extents of other users or inactive system extents. The security level and discretionary access permissions on the devices shall determine who may mount any extent on that device. The disk devices shall not be reserved only for KSOS file systems. A given disk device may be assigned for private, non-file system use.

The following checks shall be performed prior to performing the mount:

- a. the process must be privileged to issue the call
- b. the user must own the device
- c. the user must have write permission to the extent
- d. the security level of the device must be equal to the security level of the extent
- e. both the root SEID and the leaf SEID must be files (although, normally they would both be of the subtype directory)
- f. the leaf SEID must be inactive (not opened)
- g. the device SEID must be a block organized device (disk)
- h. the file system must not be presently mounted
- i. sufficient Kernel mount table space must exist (this is a low bandwidth confinement channel)

The Kernel primitive K_mount would normally be called by privileged Non-Kernel Security-Related Software which would make additional checks before issuing the K_mount call.

Security Kernel B5 Specification

3.2.13.3 Outputs

The Kernel primitive `K_mount` shall return the following information.

ec: error conditions

The Kernel primitive `K_mount` shall detect and return the following exception (error) conditions.

- EX: insufficient privilege level
- EX: the leaf file does not exist or is inaccessible for security or integrity reasons
- EX: the extent, dev, does not exist or is inaccessible for security or integrity reasons
- EX: the leaf file cannot be written by this user (discretionary access)
- EX: the dev extent cannot be written by this user
- EX: the leaf is not a file
- EX: the root is not a file
- EX: dev is not an extent
- EX: the extent is already mounted
- EX: this user is not the owner of the extent
- EX: the extent is open
- EX: the leaf is open
- EX: no more room in the mount table

3.2.14 Kernel Function: K_nap

3.2.14.1 Inputs

The Kernel primitive `K_nap` shall take the following parameters.

timeOut: time which should pass before process should be rescheduled

3.2.14.2 Processing

`K_nap` shall be a mechanism for explicitly giving up the processor when a higher level blocking condition occurs. This situation occurs when, for example, processes implementing semaphores on top of the Kernel become logically blocked on a semaphore. `K_nap` provides an alternative to busy waiting for the semaphore. The timeOut argument shall be the incremental time before which the process should not be rescheduled by the Kernel. Processes using `K_nap` should check that the logical condition for which they were waiting has occurred when they are activated.

3.2.14.3 Outputs

The Kernel primitive `K_nap` shall return the following information.

none

The Kernel primitive `K_nap` shall detect and return the following exception

Security Kernel B5 Specification

(error) conditions.

none

3.2.15 Kernel Function: K_open

3.2.15.1 Inputs

The Kernel primitive K_open shall take the following parameters.

oSeid: SEID of the object to be opened

om: requested access mode (read, write, read and write, write with exclusive use, or read and write with exclusive use)

stCap: explicit subtype capability (optional, open object descriptor)

3.2.15.2 Processing

The Kernel open primitive shall be the Kernel analog of the UNIX user interface open call. SEID interpretation shall be performed by the Kernel. A successful K_open call may be thought of as being equivalent to successfully granting the requested access capability to the object. The security and integrity levels of the calling process must permit the flow of information to or from the object to be opened, depending upon the open mode. For example, to open for reading, the security level of the object must be at or below the level of the process. To open for writing, the security level of the object must be at or above the level of the process. The caller shall have the appropriate explicit or implicit discretionary permission to the subtype of the file. Explicit subtype capability shall be passed to the Kernel in the form of an open object descriptor returned from a previous, successful K_open on the file subtype SEID. Implicit subtype capability shall be derived by checking the requested access mode against the valid access modes permitted to the effective group and user ID of the process. That is, if the subtype discretionary access permissions allow accesses in the requested mode, no explicit subtype capability need be supplied. For example, if the subtype discretionary access permissions allow read access by all users and the user is opening the object for reading, the subtype capability need not be supplied. The caller shall have the appropriate discretionary access permission to the file as granted by the owner of the file. An open object descriptor (virtualized on a per process basis) shall be returned for subsequent references to the object.

The valid combinations of access modes shall be:

- a. read
- b. write
- c. read and write
- d. write with exclusive use

Security Kernel B5 Specification

e. read and write with exclusive use

As indicated above, if exclusive use is requested, write access must also be requested. This requirement is motivated by security considerations.

Requests for open with exclusive use shall be blocking. (Failure returns from unsuccessful exclusive use opens create clandestine information channels.) They shall remain blocked until the outstanding open on the object has been closed or until a (system default) timeout expires. Non-exclusive use opens shall also block when attempting to open a file with outstanding exclusive use access. The number of files opened for exclusive use shall be limited to one per process to prevent a class of deadlocks. Exclusive use opens shall not be passed to the child process as a result of a fork operation. The fork shall fail until the exclusive use file has been closed.

3.2.15.3 Outputs

The Kernel primitive `K_open` shall return the following information.

ec: error condition
od: open descriptor (process local)

The Kernel primitive `K_open` shall detect and return the following exception (error) conditions.

EX: object does not exist, or is inaccessible from security or integrity considerations
EX: object link count is zero
EX: subtype capability that was supplied was not for a subtype
EX: subtype capability not supplied, but object was in a subtype
EX: subtype capability supplied was not opened in the modes that are being requested for the object, or was not the same subtype as the object
EX: no discretionary access permission
EX: device on which object resides has been mounted as read only, and write access requested
EX: exclusive use requested, but no write access requested
EX: another process has the object open for exclusive use
EX: exclusive use requested and another file already opened for exclusive use by this process
EX: too many open objects for this process

3.2.16 Kernel Function: K_post

3.2.16.1 Inputs

The Kernel primitive `K_post` shall take the following parameters.

receiver: SEID of process to which message is to be sent
pseudoInt: Boolean flag TRUE => cause pseudo interrupt in receiving process
FALSE => no pseudo interrupt
msg: IPC message text (the exact size of the text, and whether this text is of fixed or variable size is left to the discretion of the development

Security Kernel B5 Specification

contractor, subject to Government approval)

3.2.16.2 Processing

K_post shall send a short message to another process. A pseudo interrupt shall be asserted at the destination process, if selected, and if the receiving process has IPC pseudo interrupts enabled (i.e. that its pseudo-interrupt level is sufficiently low to allow pseudo interrupts). A header shall be attached to the message indicating the SEID of the originating process.

The type independent information for the two processes shall be used to determine rights of the originating process to communicate with the the destination.

3.2.16.3 Outputs

The Kernel primitive K_post shall return the following information.

ec: error conditions

The Kernel primitive K_post shall detect and return the following exception (error) conditions.

EX: destination process does not exist or is inaccessible

EX: destination process cannot accept any more messages

Because the last exception condition can be used as an information channel, the bandwidth of this channel shall be limited by forcing a wait before return.

3.2.17 Kernel Function: K_read_block

3.2.17.1 Inputs

The Kernel primitive K_read_block shall take the following parameters.

od: open object descriptor

blockNo: the logical block number in the file at which to start reading

duFile.vloc.domain: domain of block into which the file will be read

duFile.vloc.idSpace: space of block

duFile.vloc.vAddr: starting address of the block

duFile.size: length (in bytes) of the block, i.e., the maximum number of bytes to read

as: asynchronous call identifier

3.2.17.2 Processing

K_read_block shall retrieve block(s) of data from a previously opened file. An error shall be reported if the file was not opened for reading. The security conditions and discretionary access rights shall be checked at the time of K_open. Error and end of file conditions shall be encoded into the return value from K_read_block. Other errors shall include bad parameter values and physical

Security Kernel B5 Specification

I/O errors.

Only block mode I/O shall be supported at the Kernel level. Stream I/O can be implemented in the UNIX Emulator by caching the data blocks in Emulator level buffers. The Emulator would also maintain the seek or file pointers. The block length parameter shall specify the maximum number of bytes to be read from the file or device. This parameter value may be device dependent and validity shall be checked by the specific device driver. Certain devices, such as terminals, can accept any block length. For such devices, the "block" consists of bytes from the input stream up until the occurrence of one of a user-controllable (via `K_device_function`) set of break characters or until the vector is full. The Kernel shall check for the validity of the vector in which the data from the file is to be placed. These checks include checking that the vector is valid and accessible (i.e. can be written) and that the entire vector lies within one currently mapped in segment.

The Kernel file system shall allow "sparse" files, i.e. files in which not all blocks have been written into. When one of these unwritten blocks is read zeros shall be placed in that part of the returned data vector.

Limited numbers of asynchronous `K_read_block` requests may be made. If the the asynchrony id field is non-null, `K_read_block` may return before completion of the requested operation. At the completion of the requested operation, the Kernel shall transmit an IPC message to the requesting process containing (at least): the SEID of device, the asynchrony id, the number of bytes read, and other device dependent status information.

3.2.17.3 Outputs

The Kernel primitive `K_read_block` shall return the following information.

- ac: error conditions
- ef: end of file. One of the blocks requested was larger than the number of blocks in the file. This condition may be treated as an error condition. All blocks that exist will be returned.
- da: vector of data in the block (not an explicit return parameter in the formal specifications in Appendix A)
- result.bytesRead: actual number of bytes read
- result.errst.devIndep: device independent I/O completion status
- result.erst.devDep: device dependent I/O completion status

The Kernel primitive `K_read_block` shall detect and return the following exception (error) conditions.

- EX: invalid open object descriptor
- EX: object is a subtype
- EX: object not opened for reading
- EX: object is a terminal, and this path is not the current path
- EX: size of the request was bad (i.e. below minimum, above maximum, or not an integral multiple of device dependent block size)
- EX: bad block number
- EX: end of file
- EX: bad parameter block specification (entire parameter block must fit

Security Kernel B5 Specification

within one mapped in segment which can be written into)

3.2.18 Kernel Function: K_receive

3.2.18.1 Inputs

The Kernel primitive K_receive shall take the following parameters.

timeOut: if no messages have been received within timeOut, return

3.2.18.2 Processing

K_receive shall suspend the execution of a process until the receipt of an IPC message or until a time out. The return value shall indicate the condition which caused the process to be rescheduled.

The first message in the queue of received IPC messages shall be returned. If the timeOut expires before any pseudo-interrupts occur, the error code shall so indicate.

3.2.18.3 Outputs

The Kernel primitive K_receive shall return the following information.

ec: error conditions

msg.sender: process SEID of sender (filled in by the Kernel)

msg.text: data from sender

The Kernel primitive K_receive shall detect and return the following exception (error) conditions.

EX: timeOut expired with no message received

3.2.19 Kernel Function: K_release_process

3.2.19.1 Inputs

The Kernel primitive K_release_process shall take the following parameters.

rSeid: SEID of process to release (defaults to self)

3.2.19.2 Processing

K_release_process shall deallocate all of the Kernel level resources associated with the named process. Only a process with the same owner or a process privileged to change its owner may issue a K_release_process for another process. The effects of K_close for all open files and of a K_release_segment for all the segments of the process shall occur. Shared segments shall remain intact unless the reference count to the segment has reached zero. Segments

Security Kernel B5 Specification

with a zero reference count shall be deallocated unless they have been created to be "sticky". Files shall be deallocated if their reference counts and open counts are zero. The process id shall become unknown.

3.2.19.3 Outputs

The Kernel primitive `K_release_process` shall return the following information.

ec: error conditions

The Kernel primitive `K_release_process` shall detect and return the following exception (error) conditions.

EX: not the owner of `rSeid` and no privilege to change ownership

3.2.20 Kernel Function: K release segment

3.2.20.1 Inputs

The Kernel primitive `K_release_segment` shall take the following parameters.

segDes: segment name (process local designator)

3.2.20.2 Processing

The primitive `K_release_segment` shall release the Kernel level resources associated with the specified segment. The segment shall not be deleted if other processes are still using it or if its sticky bit is set. An error condition shall be returned if the segment name is unknown within the process.

3.2.20.3 Outputs

The Kernel primitive `K_release_segment` shall return the following information.

ec: error conditions

The Kernel primitive `K_release_segment` shall detect and return the following exception (error) conditions.

EX: invalid segment designator

3.2.21 Kernel Function: K remap

3.2.21.1 Inputs

The Kernel primitive `K_remap` shall take the following parameters.

in: designator name of segment to be mapped in (may be null)

Security Kernel B5 Specification

vl.domain: domain new segment to be mapped into
vl.idSpace: Instruction space of Data space segment placement
vl.vAddr: virtual address within domain of segment
da: requested access mode (read, write, read and write)
vlFlg: Boolean flag indicating whether (TRUE) or not (FALSE) virtual location of this segment should be changed in accordance with the virtual location specified in the call
daFlg: Boolean flag indicating whether or not discretionary access modes should be changed
out: designator name of segment to be mapped out (may be null)
outSize: new size of outgoing segment
osFlg: Boolean indicating whether or not the size of the segment should be changed

3.2.21.2 Processing

The K_remap primitive shall permit the process to change its segment map. The outgoing segment shall no longer be directly addressable by the process through machine instructions. The incoming segment shall become directly addressable by the process. The outgoing segment shall not be released (deleted.) However, the memory management hardware of the segment to be removed from the current mapped set may be used to satisfy the hardware requirements of the incoming segment. When a process segment is mapped into the current addressable set of segments, it shall occupy the virtual address vector defined by the arguments to K_remap. Either but not both of the segment designators may be null. If both are null the call shall have no effects. The incoming segment must fit into the virtual memory and memory management resources available after the outbound segment is unmapped. If it does not, or if any of the other error conditions occur, the call shall have no effect on the segment mapping.

If the alter virtual location flag (vlFlg) is TRUE, the incoming segment shall be mapped into the location specified as arguments to the call, and its status information adjusted to reflect this as a permanent change. Otherwise, the segment shall be mapped into the location specified in its permanent status information.

If the alter discretionary access information flag (daFlg) is TRUE, the modes in which this process will access the segment shall be checked against the permitted access modes for the segment, and if allowed, will become the access modes for the segment. This may alter the settings of memory management hardware when the segment is mapped back in.

If the alter size flag (osFlg) is TRUE, the size of the outbound segment shall be set to the value outsize. The expansion or truncation of the segment is performed at the end of segment specified by the growth attribute of the segment specified when built. Expanded parts of segments shall be filled with zeros. The size change can only be applied to segments that are not sharable.

3.2.21.3 Outputs

The Kernel primitive K_remap shall return the following information.

Security Kernel B5 Specification

ec: error conditions

The Kernel primitive `K_remap` shall detect and return the following exception (error) conditions.

- EX: segment "in" is not one of the active segments for this process
- EX: segment "out" is not mapped in
- EX: segment "in" is already mapped in
- EX: new segment discretionary access modes are for write only
- EX: attempt to change the memory space of a segment
- EX: virtual address space overlap or insufficient memory mapping hardware to satisfy request
- EX: size change requested on sharable segment
- EX: new discretionary access modes not allowed to this segment

3.2.22 Kernel Function: K_rendezvous_segment

3.2.22.1 Inputs

The Kernel primitive `K_rendezvous_segment` shall take the following parameters.

- segSeid: name (SEID) of segment with which to rendezvous
- vl.domain: domain for the segment, i.e. which domain the segment is to be mapped into
- vl.idSpace: space for the segment
- vl.vAddr: virtual address for the segment
- da: access type requested (read | write)

3.2.22.2 Processing

The Kernel call `K_rendezvous_segment` shall be the mechanism by which processes are able to share segments. If the segment requested exists and is accessible, it shall be mapped into the processes address space as requested, providing that the requested mapping information is valid. The Kernel shall check that the segment may be mapped into the process issuing the `K_rendezvous_segment` call. The checks shall include:

- a. that the segment seid is active
- b. that the segment may be shared
- c. that the security/integrity level of the process allows it to access the segment
- d. that the discretionary access for the segment allows it to be accessed in the requested way
- e. that the virtual address supplied is valid

3.2.22.3 Outputs

The Kernel primitive `K_rendezvous_segment` shall return the following information.

ec: error conditions
sgd: process local name for the segment

The Kernel primitive `K_rendezvous_segment` shall detect and return the following exception (error) conditions.

EX: segment not active or cannot be accessed from the process security and integrity level
EX: requested access mode was write only
EX: discretionary access denied
EX: segment already known to this process (may or may not be mapped in)
EX: virtual memory address space would be exceeded or overlapped
EX: too many segments for this process

3.2.23 Kernel Function: K_secure_terminal_lock

3.2.23.1 Inputs

The Kernel primitive `K_secure_terminal_lock` shall take the following parameters.

devSeid: SEID of terminal ("desired" path)

3.2.23.2 Processing

The `K_secure_terminal_lock` call shall provide a mechanism to control which of several paths to a terminal is currently active. The call may be successfully invoked only by processes privileged to do so. For each terminal, there are several logical devices, one of which is normally reserved for use by trusted software (the "secure device"). The input/output stream is switched between these logical devices by the Secure Server (part of the Non-Kernel Security Related Software). The user may send the secure attention character to the computer. This causes the Kernel to send an IPC message to the distinguished process, the Secure Initiator [NKSIR 1978], and to "freeze" the terminal by terminating all in-progress operations on the terminal and refusing to accept new requests on any path to the terminal. Freezing means that all paths are blocked; no characters may be read or written to (from) the terminal by the active, untrusted processes. The secure attention character shall always be active, and shall be site selectable as a system generation parameter. The requirement that the secure attention character always be active means that the setting of terminal speed and parity must be performed by trusted software.

In the frozen state, no path is active, but the Secure Initiator has been informed of the situation. No further action can take place on the terminal until the trusted software causes a `K_secure_terminal_lock` to take place, thus "thawing" a path for the use of the Secure Server.

Security Kernel B5 Specification

Now, only the secure path is active, thus the user can be assured that he (she) is communicating only with trusted software, and not some untrusted, masquerading process. Only the trusted software may thaw a user path (call `K_secure_terminal_lock` with a SEID of a user path to the terminal) which then deactivates the secure path, and unblocks a normal path.

Each path to a terminal shall have a separate SEID and separate set of Type Independent Information. The SEIDs of a set of paths to a given terminal are dense; the first SEID of a set is designated the secure path. The Secure Server may change the level of a path by the use of the `K_set_object_level` function, but only if the path is currently not open by any process. Thus the tranquility property is enforced for all paths, and it is thus not possible for the level of a path to change while it is associated with a process.

A unique exception status shall be returned when a process attempts I/O on a frozen path. It is the responsibility of the process to wait a short period and try again if it wishes to wait for the path to be unfrozen (this occurs when the user changes level at his terminal and wants to leave a running process behind).

Because the processes are not killed when they are unable to write to the terminal at its new level, it is possible to return to level later and resume as if nothing had happened.

3.2.23.3 Outputs

The Kernel primitive `K_secure_terminal_lock` shall return the following information.

ec: error conditions

The Kernel primitive `K_secure_terminal_lock` shall detect and return the following exception (error) conditions.

EX: SEID does not refer to a terminal

EX: process is not privileged to use secure terminal interface

EX: SEID does not exist

3.2.24 Kernel Function: K set file status

3.2.24.1 Inputs

The Kernel primitive `K_set_file_status` shall take the following parameters.

fseid: seid of desired object

newStatus.nBlocks: block size of the file (should not be modified without corresponding alteration of any on-disk fill block allocation information)

newStatus.linkCount: number of references to the file (used by UNIX Emulator and Directory Manager to indicate number of directory entries for the file)

newStatus.timeLastMod: time of the last modification to the file

Security Kernel B5 Specification

newStatus.subtype: the subtype (if any) to which the file belongs
newStatus.openAtCrash: Boolean flag indicating whether (TRUE) or not (FALSE)
the file is open for writing or was open when the system halted (used
for recovery of higher level constructs like directories)

3.2.24.2 Processing

K_set_file_status shall set the file object type dependent information. This call should be used primarily for maintenance functions. The intent is to allow programs performing recovery and maintenance functions for higher level constructs, e.g. UNIX directories to have the ability to manipulate file status information. Maintenance and recovery of the basic Kernel file system cannot be performed completely through this call. Rather, these basic file system recovery actions may require manipulation of the disk data directly. Naturally, such programs must be designed and implemented with the same thoroughness and techniques used for the Kernel implementation, and their use must be restricted to authorized maintenance personnel. The process shall require the privilege to set file status to successfully use this call.

3.2.24.3 Outputs

The Kernel primitive K_set_file_status shall return the following information.

ec: error conditions

The Kernel primitive K_set_file_status shall detect and return the following exception (error) conditions.

EX: either file does not exist or is inaccessible from security or integrity considerations

EX: no discretionary access permissions to file

EX: not the file owner

EX: not privileged to change file control information

3.2.25 Kernel Function: K set object level

3.2.25.1 Inputs

The Kernel primitive K_set_object_level shall take the following parameters.

objseid: seid of desired object

level.nd.securityCat: new security classification category of object

level.nd.securityComp: new security compartment set

level.nd.integrityCat: new integrity classification category of object

level.nd.integrityComp: new integrity compartment set (now always null, may be omitted)

level.da: new discretionary access permissions for object (read, write, execute for owner, group and others)

level.owner: new owner for the object

Security Kernel B5 Specification

level.group: new group of the object
level.priv: new privileges for the object

3.2.25.2 Processing

The K_set_object_level primitive shall set the type independent information for an object.

The owner of the file shall be capable of changing the discretionary access information. Processes with the privilege to set object level (see 3.1.1.2.5) shall be capable of changing:

- a. the user which owns the object
- b. the group which owns the object
- c. the discretionary permissions
- d. the security level (security category and compartments)
- e. the integrity level (integrity category and (presently null) compartments.)

Processes shall require an additional privilege to set the privileges of an object. If the appropriate conditions are not met to change the object level/access information, only the information which the user is privileged to change shall be altered.

3.2.25.3 Outputs

The Kernel primitive K_set_object_level shall return the following information.

ec: error conditions

The Kernel primitive K_set_object_level shall detect and return the following exception (error) conditions.

- EX: object does not exist or is inaccessible due to security or integrity considerations
- EX: insufficient privilege to change mandatory access control information
- EX: only owner or privileged process can change file owner or discretionary permissions
- EX: improper security level, integrity level, access mode, owner or domain

3.2.26 Kernel Function: K set process status

3.2.26.1 Inputs

The Kernel primitive K_set_process_status shall take the following parameters.

Security Kernel B5 Specification

procSeid: SEID of process whose status is to be set
ps.self: the SEID of the process (cannot be altered, included for symmetry with K_get_process_status)
ps.parent: the SEID of the parent process (cannot be altered, included for symmetry)
ps.family: the process family identifier
ps.realUser: the real user ID
ps.realGroup: the real group ID
ps.pc: program counter (PDP-11 specific)
ps.ps: processor status word (PDP-11 specific)
ps.pil: pseudo-interrupt level
ps.advPrio: advisory priority
ps.timerAlarm: the time remaining before a timer interrupt will occur
ps.supervisorTiming: execution time in the supervisor domain (PDP-11 specific) (This need not be an exact timing.) (cannot be altered, included for symmetry)
ps.userTiming: execution time in the user domain (PDP-11 specific) (cannot be altered, included for symmetry)
ps.timTog: Boolean indicating whether repetitive timer interrupts are to be generated

3.2.26.2 Processing

The K_set_process_status call shall permit the process to change those type dependent parameters that are not controlled by other primitives.

The K_set_process_status Kernel call shall supply an advisory scheduling priority to the Kernel level scheduler. The Kernel may elect to adjust the advisory priority to guarantee equitable service to all processes. Only processes privileged to set privileges may alter their privileges.

The notion of real and effective user identification shall be retained at the Kernel level because these identifiers determine the access permissions extended to a process. The effective user and group ID's are part of the type independent information for the process, because they are what determine the discretionary access rights. The real user and group ID's are part of this type dependent information.

The timer toggle and pseudo interrupt level control the pseudo interrupt mechanism. If the timer toggle is TRUE, a pseudo interrupt shall be generated every clock tick (machine dependent time unit). This mechanism may be used for periodic sampling of user mode program counter values for the construction of execution profiles. The pseudo interrupt level is analogous to the hardware interrupt level. Pseudo-interrupts shall be transmitted to the process only if the level of the pseudo interrupt is above the level of the process. The process shall not be able to raise the pseudo interrupt level sufficiently to prevent hardware error pseudo interrupts.

3.2.26.3 Outputs

The Kernel primitive K_set_process_status shall return the following information.

Security Kernel B5 Specification

ec: error conditions

The Kernel primitive `K_set_process_status` shall detect and return the following exception (error) conditions.

EX: the process does not exist, or is inaccessible due to security, integrity or discretionary access considerations

3.2.27 Kernel Function: `K_set_segment_status`

3.2.27.1 Inputs

The Kernel primitive `K_set_segment_status` shall take the following parameters.

segSeid: seid of the desired segment
st.gl.sharable: Boolean indicating whether or not the segment may be shared
st.gl.lock: Boolean indicating whether or not the segment is to be locked in memory (requires privilege to set)
st.gl.sticky: Boolean indicating whether or not the segment is to be retained in the swap area after there are no more active references to it (requires privilege to set)
st.gl.memAdvise: Boolean indicating whether or not it is expected that I/O will be done to this segment (this is merely advisory, I/O can be done to any segment which is accessible in the desired mode)
st.gl.executable: Boolean indicating whether or not the segment contains executable data
st.gl.direction: growth direction for the segment (Up or down) (PDP-11 specific)
st.size: the size of the segment, in the PDP-11 this is in bytes

3.2.27.2 Processing

The Kernel primitive `K_set_segment_status` supports modification of the type dependent information of a segment. The invoking process shall have appropriate privilege in order to modify the "sticky" flag or the "lock" flag.

3.2.27.3 Outputs

The Kernel primitive `K_set_segment_status` shall return the following information.

ec: error conditions

The Kernel primitive `K_set_segment_status` shall detect and return the following exception (error) conditions.

EX: segment does not exist or is inaccessible from security or integrity considerations
EX: discretionary access violation (user must have read and write access)
EX: segment is sharable

Security Kernel B5 Specification

EX: growth direction is different than current value
EX: attempt to alter "sticky" without privilege to do so
EX: attempt to alter "lock" without privilege to do so

3.2.28 Kernel Function: K signal

3.2.28.1 Inputs

The Kernel primitive K_signal shall take the following parameters.

procSeid: SEID of process to be signaled
sigVal: value of the signal

3.2.28.2 Processing

The K_signal primitive shall provide a means for privileged processes to transmit a high priority pseudo-interrupt to a process. K_signal differs from the K_post IPC mechanism in several ways. First, K_signal always generates a pseudo interrupt. The pseudo-interrupt level of the K_signal pseudo-interrupt is above that of normal IPC. Second, the K_signal pseudo interrupt will abort long running Kernel calls (i.e. terminal I/O) which receiving the K_post mechanism does not. Third, a much smaller message is transmitted. The intended use of K_signal is to provide a mechanism for a privileged process to "get through" to another process, typically to ask it to terminate.

3.2.28.3 Outputs

The Kernel primitive K_signal shall return the following information.

ec: error conditions

The Kernel primitive K_signal shall detect and return the following exception (error) conditions.

EX: process does not have the privilege to issue the call
EX: either the recipient does not exist or is inaccessible due to security or integrity considerations
EX: recipient pseudo interrupt level too high

3.2.29 Kernel Function: K spawn

3.2.29.1 Inputs

The Kernel primitive K_spawn shall take the following parameters.

immSeid: rendezvous name or segment seid of the intermediary process segment
arg: segment designator specifying argument segment

3.2.29.2 Processing

The Kernel primitive `K_spawn` shall combine the functions of `K_fork` and `K_invoke` into one operation. The `K_spawn` primitive permits process creation without the cost of copying the parent process image to the child process. The effect of `K_spawn` shall be to create a new process and to force the effect of a `K_invoke` call upon the newly created process. The parent process may therefore completely specify the contents of the child process image.

The parameters to `K_spawn` shall be the same as the parameters to the `K_invoke` primitive. These parameters shall be used to determine the effect of the `K_invoke` call forced upon the child process. (See `K_invoke` above for a discussion of this primitive.) The full semantics of `K_invoke` shall be implemented. Hence, a child process may acquire more privilege than the parent and may operate in a different discretionary access domain.

3.2.29.3 Outputs

The Kernel primitive `K_spawn` shall return the following information.

childSeid: SEID process identifier of the child process (returned to parent, child process is returned the process SEID of the parent process)
ownSeid: process SEID (parent gets parentSeid, child gets childSeid)
ec: error conditions

The Kernel primitive `K_spawn` shall detect and return the following exception (error) conditions.

EX: unable to create new process (See note in "`K_fork`" above.)
EX: invalid argument segment designator
EX: intermediary segment does not exist or is inaccessible due to security or integrity considerations
EX: discretionary access violation on intermediary segment
EX: argument segment is not writable
EX: argument segment sharable
EX: argument segment does not fit in space allowed
EX: intermediary segment does not fit in space allowed

3.2.30 Kernel Function: K special function

3.2.30.1 Inputs

The Kernel primitive `K_special_function` shall take the following parameters.

fn: special function requested
parm.domain: domain of parameter block
parm.idSpace: space of parameter block
parm.vAddr: address of parameter block
parm.size: size of parameter block (may be omitted and default to a standard size)

3.2.30.2 Processing

The K_special_function primitive shall provide a convenient entry point to request infrequently executed functions. At the discretion of the development contractor, subject to Government approval, this primitive may be omitted, and explicit calls provided for each function required. The major functions assigned to K_special_function are:

a. shutdown processing

Processes requesting shutdown or to set the system high security and integrity levels shall have at least OPERATOR integrity level. The virtual address of the parameter block shall be checked. The entire parameter block shall lie within one segment.

3.2.30.3 Outputs

The Kernel primitive K_special_function shall return the following information.

ec: error conditions

The Kernel primitive K_special_function shall detect and return the following exception (error) conditions.

EX: insufficient integrity level for function requested

EX: invalid parameter block

3.2.31 Kernel Function: K_unlink

3.2.31.1 Inputs

The Kernel primitive K_unlink shall take the following parameters.

fSeid: SEID of file to be unlinked

3.2.31.2 Processing

K_unlink shall decrement the file reference count of the specified file. When the file reference count has decreased to zero, the file shall be deleted after any processes having it open either close it or terminate. If the reference count is zero but the file has not been deleted due to processes having it open, it cannot be opened. The security and integrity checking shall be as if the file were being opened for reading and writing. No discretionary access checking shall be done by the Kernel, allowing processes privileged to use this primitive to perform whatever checking they chose to.

The use of K_unlink may be optionally restricted to processes privileged to issue the call. This option should be exercised when the UNIX Emulator (or similar subsystems including higher level directories) will be included in the system. The unhindered use of K_unlink can harm the integrity of the UNIX

Security Kernel B5 Specification

directory structure created by the Emulator. In cases in which no Emulator will be included, the restriction on the use of K_unlink may be lifted.

The Directory Manager of the UNIX Emulator [Emulator 78] should perform additional access checks in the emulation of the UNIX unlink(II) call. In particular, the user should have write access to the directory from which the name is being deleted.

3.2.31.3 Outputs

The Kernel primitive K_unlink shall return the following information.

ec: error conditions

The Kernel primitive K_unlink shall detect and return the following exception (error) conditions.

- EX: insufficient privilege (optional, see above)
- EX: object does not exist or is inaccessible due to security or integrity considerations
- EX: object was not a file (devices cannot be unlinked, their SEIDs are permanently assigned)
- EX: object resides on a read-only file system

3.2.32 Kernel Function: K_unmount

3.2.32.1 Inputs

The Kernel primitive K_unmount shall take the following parameters.

devSeid: device SEID of file system to be dismounted

3.2.32.2 Processing

The Kernel primitive K_unmount shall logically unmount the file system specified by the device SEID. To unmount a file system the following checks must be satisfied:

- a. the process must have the privilege to issue the call
- b. the user must own the device
- c. the device must have file system mounted on it
- d. the user must have write permission to the extent
- e. the extent must be quiescent (no open files)

Security Kernel B5 Specification

3.2.32.3 Outputs

The Kernel primitive `K_unmount` shall return the following information.

ec: error conditions

The Kernel primitive `K_unmount` shall detect and return the following exception (error) conditions.

EX: insufficient privilege

EX: device not mounted or is inaccessible due to security or integrity considerations

EX: process owner (user ID) is not the same as the device owner (user ID)

EX: process does not have write discretionary premissions to the device

EX: files still open on the device

3.2.33 Kernel Function: K_walk_process_table

3.2.33.1 Inputs

The Kernel primitive `K_walk_process_table` shall take the following parameters.

n: global process table entry requested

3.2.33.2 Processing

The `K_walk_process_table` primitive shall be a means for privileged software to obtain the SEIDs of active processes. The primitive shall return the SEID of the process which occupies slot n of the global process table. This SEID can then be used in `K_get_level` or `K_get_process_status` calls. The call shall fail if the process does not possess the privilege to issue it, or if n is not a valid index number for the process table.

3.2.33.3 Outputs

The Kernel primitive `K_walk_process_table` shall return the following information.

ec: error conditions

rSeid: SEID of the n-th entry in the process table

The Kernel primitive `K_walk_process_table` shall detect and return the following exception (error) conditions.

EX: not privileged to issue call

EX: n is not a valid process table index

Security Kernel B5 Specification

3.2.34 Kernel Function: K_write_block

3.2.34.1 Inputs

The Kernel primitive K_write_block shall take the following parameters.

od: open object descriptor
blockNo: logical block number of the file at which to start writing
duFile.vl.domain: domain in which data to write resides
duFile.vl.vAddr: address at which data starts
duFile.vl.idSpace: space of data area
duFile.size: length (in bytes) of data to write
id: asynchronous call identifier

3.2.34.2 Processing

K_write_block shall move the data in the area described by the parameter block into the specified block(s) of the file designated by od. Security and discretionary access conditions shall be checked at time of K_open on the file. If the file was not opened for writing (or for reading and writing) an exception shall be returned. Other exceptions shall include bad parameter values and physical I/O errors.

The Kernel shall not maintain information about where in the file the last I/O was directed. Thus, it is the responsibility of software using the Kernel to keep such information.

Limited numbers of asynchronous K_write_block requests may be made. If the the asynchronous call identifier field is non-null, K_write_block may return before completion of the requested operation. At the completion of the requested operation, the Kernel shall transmit an message to the requesting process containing (at least): the asynchronous call identifier, the number of bytes written, and other device dependent status information.

3.2.34.3 Outputs

The Kernel primitive K_write_block shall return the following information.

ec: error conditions

The Kernel primitive K_write_block shall detect and return the following exception (error) conditions.

EX: parameter block invalid, or segment designated is not readable
EX: open descriptor invalid
EX: open descriptor not opened for writing
EX: attempt to write to a blocked terminal
EX: size of block does not match device block size or is outside allowable limits for device
EX: block number out of range
EX: size of data area would result in a block number out of range.

3.2.35 Kernel Resource Manager

The KSOS Security Kernel shall include a system resource manager that is logically separate from the physical scheduling mechanisms of the Kernel. The Kernel Resource Manager may either be made part of the Kernel or may exist as an external process. The Kernel Resource Manager shall perform higher level management of:

- a. processor or CPU resources
- b. physical memory
- c. process segment swapping space

The Kernel Resource Manager shall attempt to maximize the utilization of all system-wide resources. The Resource Manager shall attempt to guarantee fair and equitable service to all processes running in the KSOS environment. Denial of service conditions shall be avoided whenever possible. Deadlock conditions shall also be avoided.

The operation of the Kernel Resource Manager requires a tightly coupled interface with the rest of the Kernel. If the Kernel Resource Manager is implemented as an external process, care shall be taken to make the interface as efficient and secure as possible.

3.2.36 Secure Terminal Interface

To assure non-spoofable communications between the user terminal and KSOS trusted software, the Security Kernel shall provide a secure terminal interface. This function shall operate as follows. One of the characters which can be typed by the user shall be reserved for the Kernel attention character. When the user transmits this character, the Kernel shall suspend all normal input/output between the terminal and the users' processes. The Kernel shall send an IPC message to the Secure Initiator [NKSR 78] notifying this process that a user at the specified terminal has requested a special secure service. Further communication with the terminal shall only be permitted to processes permitted to open the terminal using a SEID specifying secure terminal communications. Normal communications shall resume after a thaw command has been issued by an appropriately privileged process, through the `K_secure_terminal_lock` primitive.

3.2.37 Audit Messages

Standard DoD military security policy requires a record of security related operations which may be subject to audit. The KSOS Kernel shall transmit audit messages to the secure Audit Capture Process [NKSR 78]. The Security Kernel shall generate audit messages for (at least) the following security related events.

- a. file creation
- b. file deletion

Security Kernel B5 Specification

- c. any operation that changes the file owner, security level, integrity level, or process privileges associated with process image files
- d. available resource exhaustion conditions for the resources:
 - 1. available processes
 - 2. file space
 - 3. IPC message buffer space.
- e. any access attempt denied because of security or integrity violations

Audit messages may be generated for other security related events to be determined by the development contractor subject to Government review and approval. Kernel-level audit functions shall be capable of being selectively disabled during system generation at the discretion of the System Security Officer. The exact format of the audit messages and the mechanism by which they are transmitted shall be left to the development contractor subject to Government approval.

3.2.38 Kernel Trap Management

3.2.38.1 Inputs

The inputs to the Security Kernel trap management function shall be the various processor exception conditions and device interrupts. On the PDP-11/70, these conditions are vectored to specific memory locations which shall be accessible only from kernel domain software.

3.2.38.2 Processing

The Security Kernel shall dispatch device interrupts to the appropriate peripheral device handler. In the PDP-11/70 implementation of KSOS, device interrupt handling shall be limited to the Security Kernel. In other computer architectures which can guarantee the protection and security of user level device handling, these functions could be located outside the perimeter of the Security Kernel.

Certain processor exception conditions shall be reflected to the supervisor domain of a user process, e.g., the UNIX Emulator. The processor exception conditions reflected outwards shall be those traps caused by the software executing outside the domain of the Kernel. The PDP-11/70 version of KSOS, the following traps shall be reflected to supervisor domain software by means of the pseudo-interrupt mechanism:

- a. Illegal or reserved instruction exceptions. Illegal instruction codes. Reserved instructions executed in user or supervisor domain.
- b. Breakpoint traps. Execution of the BPT instruction. Trace traps.
- c. IOT instruction traps. (Causes UNIX abort.)

Security Kernel B5 Specification

- d. TRAP instruction. (Causes entry to UNIX emulator dispatch.)
- e. Floating point errors. Illegal floating point operation codes. Floating to integer conversion errors. Floating underflow and overflow. Floating undefined variable. Maintenance trap.
- f. Memory management traps. Non-resident segment trap. Segment length abort. Read only access violation. Memory management trap.

Certain trap conditions represent either system hardware failures, or serious problems internal to the Kernel. These trap conditions shall cause an unplanned system shutdown. Where possible, the system should communicate the cause of the unplanned shutdown either by sending information to the Audit Capture Process or by printing a message on the system console.

3.2.39 Pseudo Interrupts

The Kernel shall provide a pseudo interrupt facility to be used by the IPC mechanism and for trap reflection and timer interrupts. The mechanism for the PDP-11/70 implementation shall set the program counter and processor status word to values retrieved from fixed locations in the supervisor domain. The processor status word shall be masked before actually being used to prevent the user from specifying the kernel mode. Processes shall be able to inhibit pseudo interrupts through the appropriate settings of parameters in the type dependent information for the process. The `K_interrupt_return` call shall be the mechanism for resuming normal processing after a pseudo interrupt. The call shall restore the state of the process from the values contained in the pseudo interrupt vector. These locations may be modified by the software involved in interrupt processing. The Kernel shall prevent the process from increasing its accessible domains by masking the processor status word before restoring it.

3.2.40 Special Requirements

Because the Security Kernel CPCI is the heart of the system's security, it shall be produced with extraordinary precision and thoroughness. The System Specification (Type A) [A-Specs 78] describes the design and implementation standards to be followed in the Security Kernel CPCI. The Security Kernel shall be designed and implemented in a manner which is amenable to eventual formal proof of the security properties of the system.

The Security Kernel shall be designed to facilitate its modification and enhancement. It should be recognized that any modifications to the Security Kernel may invalidate all or part of the proofs of system security. Thus, modification efforts must be carefully controlled. No field maintenance of the Security Kernel shall be permitted.

3.2.40.1 Human Performance

This paragraph is not applicable to this specification.

Security Kernel B5 Specification

3.2.40.2 Government Furnished Property List

The Security Kernel CPCI shall be implemented in a Government-approved language. The compiler for this language may be furnished by the Government.

3.3 Adaptation

The Security Kernel shall be capable of adapting to any of the hardware configurations shown in the System Specification [A-Specs 78]. The creation of a Security Kernel for a particular configuration shall be an automated process designed to be performed by personnel without extensive computer knowledge. The mechanism shall employ checksums or similar techniques to assure the integrity of the inputs to the system generation process. The Security Kernel shall operate correctly in a configuration that is reduced from the one for which it was generated. This KSOS System Generation mechanism is discussed in the Non-Kernel Security Related Software CPCI Specification [NKSr 78].

3.3.1 General Environment

The Kernel and the Non-Kernel Security Related Software shall establish and enforce a maximum security level for the system as a whole and for each peripheral device and disk extent. In specifying the maximum level for peripheral devices, the System Security Officer should consider the physical location of the device, the protection of any connecting communication lines, and the personnel which would have access to the device. Individual sites may be governed by specific policies in this area.

3.3.2 System Parameters

The size of all system tables and arrays shall be parameterized to allow them to be changed easily. All manifest constants shall also be parameterized.

3.3.3 System Capacities

The Security Kernel shall be capable of supporting the maximum hardware configuration defined in the System Specifications. All terminals in that configuration shall be capable of being simultaneously on-line. The Security Kernel shall be capable of supporting at least fifty (50) active processes.

Security Kernel B5 Specification

4. QUALITY ASSURANCE PROVISIONS

4.1 Introduction

The underlying philosophy of the overall KSOS quality assurance is to provide a convincing demonstration of the security and completeness of the system. Testing and quality assurance are a complement to the formal verification aspects of KSOS. All testing on KSOS shall be conducted at the development contractor's facility.

KSOS shall be subject to the technical reviews and audit procedures of MIL-STD-1521A, Appendices A-F. The following technical reviews and audits will be conducted:

- a. System Requirements Review (Appendix A)
- b. System Design Review (Appendix B)
- c. Preliminary Design Review (Appendix C)
- d. Critical Design Review (Appendix D)
- e. Functional Configuration Audit (Appendix E)
- f. Physical Configuration Audit (Appendix F)

The Critical Design Review shall include a review of any mathematical proofs showing that the design meets the requirements of the Government approved DoD security model.

The Functional Configuration Audit shall be the vehicle for the formal review of the design versus the mathematical description. The Physical Configuration Audit shall be the vehicle for the formal review of the "as built" computer programs versus both their design and their supporting documentation. In addition, the Physical Configuration Audit shall verify that the implementation requirements of the Type A System Specifications have been met. Both audits shall include review of any relevant mathematical proofs of correctness.

4.1.1 Category I Testing

The demonstration that the requirements of Section 3 have been met shall consist of a combination of formal inspection of the CPCI and tests which demonstrate that the CPCI meets its functional requirements. The tests on the Security Kernel shall exercise each of the options for each primitive. To the extent feasible, the tests shall attempt to exercise combinations of parameters to the Kernel calls. Test sequences shall include cases which are expected to fail, for example tests which deliberately attempt to violate security. All the exception conditions defined in Section 3. shall be so tested.

4.1.2 Category II Testing

The KSOS system shall be subject to an overall system test (Category II) in accordance with the System Specification (Type A) and the development contract.

Security Kernel B5 Specification

This Category II test shall demonstrate that the three CPCI's which make up the KSOS system function correctly together. The Category II testing shall attempt to represent a typical user load including at least ten (10) on-line terminals and a network connection. To the extent feasible, the Category II testing shall also be automated through the use of shell scripts, and (optionally) another computer system emulating the terminal load.

Security Kernel B5 Specification

Table I - Quality Assurance

Assurance techniques:

NA	Not Applicable
V	Verification
F	Formal Specification
T	Testing
D	Demonstration
A	Analysis
I	Inspection

Applicable test types:

1	Module level testing
2	Integration testing
3	Acceptance testing

Requirement	Assurance Techniques							Test types		
	NA	V	F	T	D	A	I	1	2	3
3.1.1.1.				x	x	x	x	x	x	x
3.1.1.2.1.1.			x	x	x	x	x	x	x	x
3.1.1.2.1.2.			x	x	x	x	x	x	x	x
3.1.1.2.1.3.			x	x	x	x	x	x	x	x
3.1.1.2.1.4.			x	x	x	x	x	x	x	x
3.1.1.2.1.5.			x	x	x	x	x	x	x	x
3.1.1.2.2.				x	x	x	x	x	x	x
3.1.1.2.3.	x									
3.1.1.2.4.				x	x	x	x	x	x	x
3.1.1.2.5.			x	x	x	x	x	x	x	x
3.1.1.2.6.				x	x	x	x	x	x	x
3.2.1. - 3.2.34		*	x	x	x	x	x	x	x	x
3.2.35.			x	x	x	x	x	x	x	x
3.2.36.				x	x	x	x	x	x	x
3.2.37.			x	x	x	x	x	x	x	x
3.2.38.			x	x	x	x	x	x	x	x
3.2.39.			x	x	x	x	x	x	x	x
3.2.40.1.	x									
3.2.40.2.	x									
3.3.						x	x			
3.3.1.						x	x			
3.3.2.							x			
3.3.3.				x	x			x	x	x

* Selected modules will be shall be formally verified on a time and materials basis.

Security Kernel B5 Specification

5. PREPARATION FOR DELIVERY

This section is not applicable to this specification.

Security Kernel B5 Specification

6. NOTES

This section is not applicable to this specification.

Security Kernel B5 Specification

10. APPENDIX

Security Kernel B5 Specification

10. APPENDIX

This appendix presents the formal specifications for the KSOS Security Kernel. The specifications describe the interface that the Kernel shall present to its users at the Kernel call level. These formal specifications and the textual descriptions of the Kernel call interface found in Section 3 are intended to be used side-by-side, each serving to illuminate the other. The formal specifications are written in the SPECIAL (SPECification and Assertion Language) language described by [Roubine and Robinson 77].

For the purposes of specification, the Kernel has been decomposed into several modules which are viewed as being hierarchically ordered. A convention has been followed that each module is referred to by a three-letter mnemonic name derived from its function.

Each of the specification modules represents an incremental abstract machine, built upon the abstract machines that come below it in the hierarchy. A module contains descriptions of the data types which it manipulates, and functions which represent the state variables (V-functions) and primitive operations (O-functions) of the abstract machine. The functions of a module may be thought of as being the mechanism of that module's abstract machine.

Not all the functions of a module are necessarily "visible to" (i.e. available for use by) other modules. Furthermore, it is a strict rule of the design methodology employed that no function at a given level in the hierarchy is ever visible below that level.

The modules contain varying amounts of mechanism. This is dictated by the extent to which the Kernel calls they mechanize alter the state of the Kernel. A ground rule followed in writing these specifications has been to specify nothing which is not somehow visible at the Kernel call interface.

All the Kernel calls, data types, and data structures which are visible at the Kernel call interface have been collected, for ease of reference into the topmost layer of the hierarchy ("KER"). The effects of any Kernel call may be determined from the specifications by finding the entry for that call in KER, and then following the flow of implication through the hierarchy.

Another useful way in which the specifications may be studied is to determine the mechanisms and structures supporting a particular Kernel feature. This is best done by examination of the module(s) which specify the feature.

SRI International has developed and maintains tools which can verify the internal consistency of a set of formal specifications written in SPECIAL. These specifications have been processed through these tools and have been judged by the tools to be syntactically correct and semantically consistent. In the future these specifications will be processed by other tools to verify that they do indeed provide a system consistent with the mathematical security model.


```

1  $("      MODULE:      ker.specs (version 4.8)
2          CONTENTS:    Kernel Calls
3          TYPE:        SPECIAL.specifications
4          LAST CHANGED: 12/23/80, 16:51:22
5  ")
6
7
8  MODULE ker
9
10
11  TYPES
12
13  $(from mac)
14  vAddrType: {0 .. MACmaxVAddr};
15
16  $(from smx)
17  nonDisType: STRUCT_OF(
18      INTEGER securityLevel; SET_OF securityCat securityCatS;
19      INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
20  daType: SET_OF daMode;
21  modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
22  ttiStruct: STRUCT_OF(nonDisType nd;
23      modeStruct ms; INTEGER owner, group; SET_OF privType priv);
24
25  $(FROM pvm)
26  virtualLocation:
27      STRUCT_OF(domainType domain; spaceType idSpace; vAddrType vAddr);
28  globalData:
29      STRUCT_OF(BOOLEAN sharable, swappable, sticky, memAdvise, executable);
30  statusStruct:
31      STRUCT_OF(globalData gl; direction growth; INTEGER size; seid iseid;
32          segDes udes; virtualLocation vloc; daType da; BOOLEAN mapped);
33  pBlock: STRUCT_OF(virtualLocation vloc; INTEGER size);
34
35  $(FROM fca)
36  fileStatus: STRUCT_OF(INTEGER nBlocks, linkCount, timeLastMod; seid subtype;
37      BOOLEAN openAtCrash);
38  asyncId: CHAR;
39  ioStatus: STRUCT_OF(INTEGER devIndep, devDep);
40  fileNsp: {129 .. 255};      $(nsp values for file systems)
41
42  $(from pro)
43  piLevelType: {PROminPiLevel .. PROmaxPiLevel}; $(pseudo interrupt level range)
44  piEntryType: STRUCT_OF(piLevelType oldPil;
45      piLevelType oldPil;
46      INTEGER oldPc;
47      INTEGER oldPs;
48      INTEGER parameter;
49      INTEGER newPc;
50      INTEGER newPs);
51  piVectorType: {VECTOR_OF piEntryType piv | LENGTH(piv) = PROmaxPiLevel - PROminPiLevel};
52  ipcqType: {VECTOR_OF ipcMessageType zz | LENGTH(zz) <= IPCmaxMessageCount};
53  ipcTextType: {VECTOR_OF CHAR vc | LENGTH(vc) = IPCmaxMessageLength};
54  ipcMessageType: STRUCT_OF(seid sender; ipcTextType text);
55  processStatusType: STRUCT_OF(seid self;
56      seid parent;

```

```

57         INTEGER family;
58         INTEGER realUser;
59         INTEGER realGroup;
60         INTEGER pc;           $(program counter)
61         INTEGER ps;           $(processor status)
62         piLevelType pil;
63         piVectorType piv;
64         ipcqType ipcq;
65         INTEGER advPrio;       $(advisory priority)
66         INTEGER timerAlarm;    $(one'zero crossing => pi)
67         INTEGER supervisorTiming;
68         INTEGER userTiming;
69         BOOLEAN timTog;        $(timer toggle TRUE is ON)
70
71
72     EXTERNALREFS
73
74     FROM mac:
75     INTEGER MACmaxVAddr;
76
77     FROM smx:
78     seid: DESIGNATOR;
79         privType: {
80             privFileUpdateStatus,  privLink,
81             privLockSeg,           privModifyPriv,
82             privMount,             privSetLevel,
83             privStickySeg,         privSetPath,
84             privViolSimpSecurity,  privViolStarSecurity,
85             privViolSimpIntegrity, privViolStarIntegrity,
86             privViolDiscrAccess,  privRealizeExecPermission,
87             privSignal,           privWalkPTable,
88             privHalt,             privKernelCall,
89             privViolCompartments, privSetComm,
90             privImmigrate
91         };
92
93     daMode: {daRead, daWrite, daExecute};
94     securityCat: DESIGNATOR;
95     integrityCat: DESIGNATOR;
96     domainType: {nullDomain, kernelDomain, supervisorDomain, userDomain};
97
98     FROM pvm:
99     segDes: DESIGNATOR;
100    spaceType: {nullSpace, iSpace, dSpace};
101    direction: {up, down};
102    OVFUN PVMbuild(seid pSeid; statusStruct ss; modeStruct ms)
103    ^> STRUCT_OF(seid segSeid; segDes segd) result;
104    OFUN PVMdestroy(seid pSeid; segDes segd);
105    VFUN PVMgetSegmentStatus(seid pSeid, segSeid; segDes segd) ^> statusStruct st;
106    OFUN PVMsetSegmentStatus(seid pSeid, segSeid; globalData gl);
107    OFUN PVMremap(seid pSeid; segDes in; virtualLocation vl; daType d;
108        BOOLEAN vlFlg, daFlg; segDes out; INTEGER outSize;
109        BOOLEAN osFlg);
110    OVFUN PVMrendezvous(seid pSeid, segSeid; virtualLocation vl; daType d)
111    ^> segDes segd;
112

```

```

113 FROM fca:
114 openDescriptor: DESIGNATOR;
115 openModes: {omRead, omWrite, omExclusiveRead, omExclusiveWrite};
116 IOfunction: {rewind, etc};
117 OFUN FCAClose(seid pSeid; openDescriptor od);
118 OVFUN FCAcreate(seid pSeid, nspSeid; SET_OF openModes mode; modeStruct d;
119   openDescriptor stCap)
120   "> STRUCT_OF(seid fSeid; openDescriptor od) result;
121 OVFUN FCAopen(seid pSeid, oSeid; SET_OF openModes om; openDescriptor stCap)
122   "> openDescriptor od;
123 VFUN FCAgetFileStatus(seid pSeid, fSeid) "> fileStatus fst;
124 OFUN FCAsetFileStatus(seid pSeid, fSeid; fileStatus newfst);
125
126 FROM fmi:
127 OFUN FCALink(seid pSeid, fSeid);
128 OVFUN FCAmount(seid pSeid, dev, root; BOOLEAN readOnly) "> fileNsp newNsp;
129 OFUN FCAterminalLock(seid pSeid, devSeid);
130 OFUN FCAunLink(seid pSeid, fSeid);
131 OFUN FCAunmount(seid pSeid; fileNsp fsNsp);
132
133
134 FROM pro:
135 INTEGER IPCmaxMessageCount;
136 INTEGER IPCmaxMessageLength;
137 INTEGER PROminPiLevel;
138 INTEGER PROmaxPiLevel;
139 OVFUN PROfork(seid pSeid) "> seid childSeid;
140 VFUN PROgetProcessStatus(seid pSeid, getSeid) "> processStatusType ps;
141 OFUN PROinterruptReturn(seid pSeid);
142 OFUN PROinvoke(seid pSeid, immSeid; segDes arg);
143 OFUN PROnap(seid pSeid; INTEGER timeOut);
144 OFUN PROpost(seid pSeid, receiver; BOOLEAN pseudoInt; ipcTextType msg);
145 OVFUN PROreceive(seid pSeid; INTEGER timeOut) "> ipcMessageType msg;
146 OFUN PROreleaseProcess(seid pSeid, rSeid);
147 OFUN PROsetProcessStatus(seid pSeid, procSeid; processStatusType ps);
148 OFUN PROsignal(seid sender, receiver; ipcTextType signalMsg);
149 OVFUN PROspawn(seid pSeid, immSeid; segDes arg) "> seid childSeid;
150 VFUN PROwalkProcessTable(seid pSeid; INTEGER n) "> seid rSeid;
151
152 FROM lev:
153 VFUN LEVgetObjectLevel(seid pSeid, objSeid) "> tiiStruct level;
154 OFUN LEVsetObjectLevel(seid pSeid, objSeid; tiiStruct level);
155
156 FROM spf:
157 SPFFunctionType: {syncSPF, immSegSPF, sysHaltSPF, levelSetSPF};
158
159 FROM pbl:
160 OFUN PBLdeviceFunction(seid pSeid; openDescriptor od; IOfunction f;
161   pBlock arguments, status; asyncId id);
162 OVFUN PBLreadBlock(seid pSeid; openDescriptor od; INTEGER blockNo;
163   pBlock duFile; asyncId as)
164   "> STRUCT_OF(INTEGER bytesRead; ioStatus errst) result;
165 OFUN PBLspecialFunction(seid pSeid; SPFFunctionType fn; pBlock parm);
166 OVFUN PBLwriteBlock(seid pSeid; openDescriptor od; INTEGER blockNo;
167   pBlock duFile; asyncId as)
168   "> ioStatus ios;

```

```

169
170
171     FUNCTIONS
172
173     $("Note regarding implementation: The implicit parameter pSeid is
174     implemented by passing a null seid to the kernel interface. The null
175     seid is converted upon entering the interface to the seid of the
176     calling process. As a consequence, for a process to learn its own
177     seid, it suffices to call K_get_Process_Status with a null seid.")
178
179     $(Visible Kernel Functions in Alphabetical Order)
180
181     OVFUN K_build_segment(statusStruct ss; modeStruct ms)
182         [seid pSeid] ^> STRUCT_OF(seid segSeid; segDes segd) result;
183                                     $(K_build_segment)
184     EXCEPTIONS
185         EXCEPTIONS_OF PVMbuild(pSeid, ss, ms);
186     EFFECTS
187         result = EFFECTS_OF PVMbuild(pSeid, ss, ms);
188
189     OFUN K_close(openDescriptor od)[seid pSeid];                                $(K_close)
190     EXCEPTIONS
191         EXCEPTIONS_OF FCAClose(pSeid, od);
192     EFFECTS
193         EFFECTS_OF FCAClose(pSeid, od);
194
195     OVFUN K_create(seid nspSeid; SET_OF openModes mode; modeStruct d;
196         openDescriptor stCap)[seid pSeid]
197         ^> STRUCT_OF(seid fSeid; openDescriptor od) return;                $(K_create)
198     EXCEPTIONS
199         EXCEPTIONS_OF FCAcreate(pSeid, nspSeid, mode,d, stCap);
200     EFFECTS
201         return = EFFECTS_OF FCAcreate(pSeid, nspSeid, mode,d, stCap);
202
203     OFUN K_device_function(openDescriptor od; IOfunction f;
204         pBlock arguments, status; asyncId id)
205         [seid pSeid];                                                    $(K_device_function)
206     EXCEPTIONS
207         EXCEPTIONS_OF PBLdeviceFunction(pSeid, od, f, arguments, status, id);
208     EFFECTS
209         EFFECTS_OF PBLdeviceFunction(pSeid, od, f, arguments, status, id);
210
211     OVFUN K_fork()[seid pSeid] ^> seid childSeid;                            $(K_fork)
212     EXCEPTIONS
213         EXCEPTIONS_OF PROfork(pSeid);
214     EFFECTS
215         childSeid = EFFECTS_OF PROfork(pSeid);
216         $("ChildSeid is returned to the (original) process; pSeid (parent seid)
217         is returned to the newly created child")
218
219     VFUN K_get_file_status(seid fSeid)[seid pSeid] ^> fileStatus fst;
220                                     $(K_get_file_status)
221     EXCEPTIONS
222         EXCEPTIONS_OF FCAgetFileStatus(pSeid, fSeid);
223     DERIVATION
224         FCAgetFileStatus(pSeid, fSeid);

```

```

225
226 VFUN K_get_object_level(seid objSeid)[seid pSeid] ^> ttiStruct ot1;
227                                     $(K_get_object_level)
228 EXCEPTIONS
229 EXCEPTIONS_OF LEVgetObjectLevel(pSeid, objSeid);
230 DERIVATION
231 LEVgetObjectLevel(pSeid, objSeid);
232
233 VFUN K_get_process_status(seid getSeid)[seid pSeid] ^> processStatusType ps;
234                                     $(K_get_process_status)
235 EXCEPTIONS
236 EXCEPTIONS_OF PROgetProcessStatus(pSeid, getSeid);
237 DERIVATION
238 PROgetProcessStatus(pSeid, getSeid);
239
240 VFUN K_get_segment_status(seid segSeid; segDes segd)[seid pSeid] ^>
241                                     statusStruct st;
242                                     $(K_get_segment_status)
243 EXCEPTIONS
244 EXCEPTIONS_OF PVMgetSegmentStatus(pSeid, segSeid, segd);
245 DERIVATION
246 PVMgetSegmentStatus(pSeid, segSeid, segd);
247
248 OFUNK K_interrupt_return()[seid pSeid];                                     $(K_interrupt_return)
249 EXCEPTIONS
250 EXCEPTIONS_OF PROinterruptReturn(pSeid);
251 EFFECTS
252 EFFECTS_OF PROinterruptReturn(pSeid);
253
254 OFUN K_invoke(seid immSeid; segDes arg)[seid pSeid];                         $(K_invoke)
255 EXCEPTIONS
256 EXCEPTIONS_OF PROinvoke(pSeid, immSeid, arg);
257 EFFECTS
258 EFFECTS_OF PROinvoke(pSeid, immSeid, arg);
259
260 OFUN K_link(seid fSeid)[seid pSeid];                                         $(K_link)
261 EXCEPTIONS
262 EXCEPTIONS_OF FCAlink(pSeid, fSeid);
263 EFFECTS
264 EFFECTS_OF FCAlink(pSeid, fSeid);
265
266 OVFUN K_mount(seid dev; BOOLEAN readOnly)[seid pSeid]                       $(K_mount)
267                                     ^> fileNsp newNsp;
268 EXCEPTIONS
269 EXCEPTIONS_OF FCAmount(pSeid, dev, readOnly);
270 EFFECTS
271 newNsp = EFFECTS_OF FCAmount(pSeid, dev, readOnly);
272
273 OFUN K_nap(INTEGER timeOut)[seid pSeid];                                     $(K_nap)
274 EXCEPTIONS
275 EXCEPTIONS_OF PROnap(pSeid, timeOut);
276 EFFECTS
277 EFFECTS_OF PROnap(pSeid, timeOut);
278
279 OVFUN K_open(seid oSeid; SET_OF openModes om; openDescriptor stCap)
280 [seid pSeid] ^> openDescriptor od;                                         $(K_open)

```

```

281 EXCEPTIONS
282 EXCEPTIONS_OF FCAopen(pSeid, oSeid, om, stCap);
283 EFFECTS
284 od = EFFECTS_OF FCAopen(pSeid, oSeid, om, stCap);
285
286 OFUN K_post(seid receiver; BOOLEAN pseudoInt; ipcTextType msg)[seid pSeid];
287 $(K_post)
288 EXCEPTIONS
289 EXCEPTIONS_OF PROpost(pSeid, receiver, pseudoInt, msg);
290 EFFECTS
291 EFFECTS_OF PROpost(pSeid, receiver, pseudoInt, msg);
292
293 OFUN K_read_block(openDescriptor od; INTEGER blockNo; pBlock duFile;
294 asyncId as)[seid pSeid]
295 *> STRUCT_OF(INTEGER bytesRead; ioStatus errst) result; $(K_read_block)
296 EXCEPTIONS
297 EXCEPTIONS_OF PBLreadBlock(pSeid, od, blockNo, duFile, as);
298 EFFECTS
299 result = EFFECTS_OF PBLreadBlock(pSeid, od, blockNo, duFile, as);
300
301 OFUN K_receive(INTEGER timeOut)[seid pSeid] *> ipcMessageType msg;
302 $(K_receive)
303 EXCEPTIONS
304 EXCEPTIONS_OF PROreceive(pSeid, timeOut);
305 EFFECTS
306 msg = EFFECTS_OF PROreceive(pSeid, timeOut);
307
308 OFUN K_release_process(seid rSeid)[seid pSeid]; $(K_release_process)
309 EXCEPTIONS
310 EXCEPTIONS_OF PROreleaseProcess(pSeid, rSeid);
311 EFFECTS
312 EFFECTS_OF PROreleaseProcess(pSeid, rSeid);
313
314 OFUN K_release_segment(segDes segd)[seid pSeid]; $(K_release_segment)
315 EXCEPTIONS
316 EXCEPTIONS_OF PVMdestroy(pSeid, segd);
317 EFFECTS
318 EFFECTS_OF PVMdestroy(pSeid, segd);
319
320 OFUN K_remap(segDes in; virtualLocation vl; daType d; BOOLEAN vlFlg, daFlg;
321 segDes out; INTEGER outSize; BOOLEAN osFlg)[seid pSeid];
322 $(K_remap)
323 EXCEPTIONS
324 EXCEPTIONS_OF
325 PVMremap(pSeid, in, vl, d, vlFlg, daFlg, out, outSize, osFlg);
326 EFFECTS
327 EFFECTS_OF
328 PVMremap(pSeid, in, vl, d, vlFlg, daFlg, out, outSize, osFlg);
329
330 OFUN K_rendezvous_segment(seid segSeid; virtualLocation vl; daType d)
331 [seid pSeid]
332 *> segDes segd; $(K_rendezvous_segment)
333 EXCEPTIONS
334 EXCEPTIONS_OF PVMrendezvous(pSeid, segSeid, vl, d);
335 EFFECTS
336 segd = EFFECTS_OF PVMrendezvous(pSeid, segSeid, vl, d);

```

```

337
338 OFUN K_secure_terminal_lock(seid devSeid)[seid pSeid];
339                                     $(K_secure_terminal_lock)
340 EXCEPTIONS
341     EXCEPTIONS_OF FCATerminalLock(pSeid, devSeid);
342 EFFECTS
343     EFFECTS_OF FCATerminalLock(pSeid, devSeid);
344
345 OFUN K_set_file_status(seid fSeid; fileStatus fst)[seid pSeid];
346                                     $(K_set_file_status)
347 EXCEPTIONS
348     EXCEPTIONS_OF FCASetFileStatus(pSeid, fSeid, fst);
349 EFFECTS
350     EFFECTS_OF FCASetFileStatus(pSeid, fSeid, fst);
351
352 OFUN K_set_object_level(seid objSeid; ttiStruct level)[seid pSeid];
353                                     $(K_set_object_level)
354 EXCEPTIONS
355     EXCEPTIONS_OF LEVsetObjectLevel(pSeid, objSeid, level);
356 EFFECTS
357     EFFECTS_OF LEVsetObjectLevel(pSeid, objSeid, level);
358
359 OFUN K_set_process_status(seid procSeid; processStatusType ps)[seid pSeid];
360                                     $(K_set_process_status)
361 EXCEPTIONS
362     EXCEPTIONS_OF PROsetProcessStatus(pSeid, procSeid, ps);
363 EFFECTS
364     EFFECTS_OF PROsetProcessStatus(pSeid, procSeid, ps);
365
366 OFUN K_set_segment_status(seid segSeid; globalData gl)[seid pSeid];
367                                     $(K_set_segment_status)
368 EXCEPTIONS
369     EXCEPTIONS_OF PVMsetSegmentStatus(pSeid, segSeid, gl);
370 EFFECTS
371     EFFECTS_OF PVMsetSegmentStatus(pSeid, segSeid, gl);
372
373 OFUN K_signal(seid procSeid; ipcTextType sigMsg)[seid pSeid];           $(K_signal)
374 EXCEPTIONS
375     EXCEPTIONS_OF PROsignal(pSeid, procSeid, sigMsg);
376 EFFECTS
377     EFFECTS_OF PROsignal(pSeid, procSeid, sigMsg);
378
379 OVFUN K_spawn(seid immSeid; segDes arg)[seid pSeid] ^> seid childSeid;
380                                     $(K_spawn)
381 EXCEPTIONS
382     EXCEPTIONS_OF PROspawn(pSeid, immSeid, arg);
383 EFFECTS
384     childSeid = EFFECTS_OF PROspawn(pSeid, immSeid, arg);
385
386 OFUN K_special_function(SPFFunctionType fn; pBlock parm)[seid pSeid];
387                                     $(K_special_function)
388 EXCEPTIONS
389     EXCEPTIONS_OF PBLspecialFunction(pSeid, fn, parm);
390 EFFECTS
391     EFFECTS_OF PBLspecialFunction(pSeid, fn, parm);
392

```

```

393 OFUN K_unlink(seid fSeid)[seid pSeid];                                $(K_unlink)
394 EXCEPTIONS
395 EXCEPTIONS_OF FCAunlink(pSeid, fSeid);
396 EFFECTS
397 EFFECTS_OF FCAunlink(pSeid, fSeid);
398
399 OFUN K_unmount(fileNsp fsNsp)[seid pSeid];                            $(K_unmount)
400 EXCEPTIONS
401 EXCEPTIONS_OF FCAunmount(pSeid, fsNsp);
402 EFFECTS
403 EFFECTS_OF FCAunmount(pSeid, fsNsp);
404
405 VFUN K_walk_process_table(INTEGER n)[seid pSeid] ^> seid rSeid;        $(K_walk_process_table)
406
407 EXCEPTIONS
408 EXCEPTIONS_OF PROwalkProcessTable(pSeid,n);
409 DERIVATION
410 PROwalkProcessTable(pSeid,n);
411
412 OVFUN K_write_block(openDescriptor od; INTEGER blockNo; pBlock duFile;
413                      asyncId id)[seid pSeid]
414 ^> ioStatus ios;                                                        $(K_write_block)
415 EXCEPTIONS
416 EXCEPTIONS_OF PBLwriteBlock(pSeid, od, blockNo, duFile, id);
417 EFFECTS
418 ios = EFFECTS_OF PBLwriteBlock(pSeid, od, blockNo, duFile, id);
419
420
421 END_MODULE

```



```

1  $("  MODULE:      ddf.specs (version 2.5)
2      CONTENTS:    Device Dependent Functions
3      TYPE:        SPECIAL.specifications
4      LAST CHANGED: 7/28/80, 09:22:51
5  ")
6
7
8  MODULE ddf
9
10 TYPES
11 deviceType: {RK05, RWPO4, RWPO5, RWPO6, RSW04, TWE16, TM11, TU56, PR11,
12             PC11, LP11, IMPI1B, LHDH};
13 deviceStruct: STRUCT_OF(BOOLEAN addressable;
14                          INTEGER minRequest, maxRequest, modSize, maxBlockNo);
15  $(properties necessary for processing IO requests for devices)
16
17 FUNCTIONS
18
19 VFUN DDFdeviceData(deviceType d) *> deviceStruct ds;          $(DDFdeviceData)
20  $(for a given type of device, defines the IO behavior)
21  INITIALLY
22    ds = (IF d = RK05
23          THEN STRUCT(TRUE, 512, 65534, 512, 203*24*1)
24          ELSE IF d INSET {RWPO4, RWPO5}
25          THEN STRUCT(TRUE, 512, 65534, 512, 22*19*411*1)
26          ELSE IF d = RWPO6
27          THEN STRUCT(TRUE, 512, 65534, 512, 22*19*411*2*1)
28          ELSE IF d = RSW04 THEN STRUCT(TRUE, 512, 65534, 512, 2048*1)
29          ELSE IF d INSET {TWE16, TM11} THEN STRUCT(FALSE, 12, 8191, 1, 0)
30          ELSE IF d = TU56 THEN STRUCT(TRUE, 512, 512, 512, 578*1)
31          ELSE IF d INSET {PR11, PC11} THEN STRUCT(FALSE, 1, 1, 1, 0)
32          ELSE IF d = LP11 THEN STRUCT(FALSE, 1, 132, 1, 0)
33          ELSE IF d INSET {IMPI1B, LHDH} THEN STRUCT(FALSE, 1, 8191, 1, 0)
34          ELSE ?);
35
36 END_MODULE

```

```

1  $("  MODULE:      fca.specs (version 4.16)
2      CONTENTS:    File Capabilities
3      TYPE:        SPECIAL.specifications
4      LAST CHANGED: 12/23/80, 16:34:45
5  ")
6
7
8  MODULE fca
9
10
11  $("This module manages all openable objects, i.e., those that are referenced
12     through the open table corresponding to a process. These objects include
13     files, devices " both addressable and nonaddressable", terminals, extents,
14     and subtypes.
15
16     Each object is identified by a seid. Seids for devices, terminals, and
17     subtypes are allocated at system generation time. These objects are
18     permanent, and cannot be dynamically allocated and deallocated.
19     Seids for files are allocated by this module. Seids for extents are
20     allocated when the device is physically mounted. Physical mounting
21     is not handled at this time " logical mounting is " but should be.
22
23     Each process at creation is assigned an open table, in which all the
24     open objects of that process are recorded, along with their mode of
25     access. The state of the open table for a process is recorded in the
26     values of the V-functions 'FCAopenTableExists(pSeid)', which tells
27     whether the open table for the process named by 'pSeid' exists, and
28     'FCAopenEntry(pSeid, od)', which gives the seid and open mode for the
29     open object of process 'pSeid' named by the open Descriptor " a
30     designator " 'od.'
31
32     The existence of an openable object is detected by a defined value for
33     'FCAfileStatusInfo(fSeid)', where 'fSeid' is the object's seid. Each
34     object's type is ascertained by looking at the nsp part of the seid.
35     Depending on the type of the object, certain V-functions hold additional
36     information. A description of this information can be found in the
37     comment directed at each type of object.")
38
39  $(" DEVICES " there are two kinds of devices, addressable and nonaddressable;
40     an addressable device, such as a disk, has two properties: it can be
41     accessed via a block number or address; and what is put onto the device
42     via a read operation is retrieved by a write when the device is read at
43     the same address. A non-addressable device, such as a tape unit, can be
44     viewed as having an infinite stream of input data and producing an
45     infinite stream of output data. Each kind of device has four quantities
46     associated with it: a minimum request, a maximum request, a size
47     modulus, and a maximum block number.
48
49     For non-addressable devices, the
50     size modulus must be 1 and the maximum block number is meaningless.
51     An IO request specifies a certain number of characters at a certain
52     block number. The number of characters must be within the range
53     defined by the minimum and maximum request quantities for the device,
54     and must be a multiple of the size modulus. The block number must be with
55     within the range {0 .. maximum block number} for the device.
56  ")

```

```

57
58 §(" TERMINALS "" With one exception, terminals are nonaddressable devices
59 whose IO requests are limited to small multiples of a single character
60 (~ 255). Terminals however, have a special property. To enable the
61 user to change the security level of his job without changing terminals,
62 there is an illusion that a single terminal is represented by a
63 multiplicity of device seids, one for each possible login security level.
64 Each seid represents a particular secure path to the terminal. Only
65 one path to the terminal may do IO operations at a time, and this path
66 is specified by the value of the V*function 'FCACurrentPath(t)', where
67 t refers to the terminal group or physical terminal.
68 The different paths associated with a particular physical terminal
69 are specified by the value of the V*function 'FCATerminalPathSet(t)',
70 where t is as above.")
71
72 §(" EXTENTS "" Extents are addressable devices, representing areas on a
73 disk, with a block size of 512 and a size determined when the device is
74 physically mounted. For the
75 purposes of this specification, the size has been predetermined, as no
76 physical mounting is specified. The special property of extents is that
77 they can be logically mounted, so that they 'become' a file system
78 or set of files. When
79 the system is started up, there are no files, except the root, only
80 extents. Each extent is mounted, setting up the actual file system
81 that a user sees when he logs on. When the extent is mounted, it can
82 no longer be accessed as an extent. Unmounting turns a file system
83 back into an extent, and the file system disappears.")
84
85 §(" FILES "" Files are addressable devices, with a block size of 512 and
86 two important properties. They can be dynamically created and deleted,
87 and they are of variable size. Writing onto the end of a file effectively
88 changes its size. Files may also be linked to. A link is a reference
89 count used by the directory manager sitting above the kernel. It
90 represents the number of directories in which a file is found. When
91 this count goes to 0, and no process has the file open, the file
92 is deleted. This is the only way of deleting files, although they can
93 be explicitly created.")
94
95 §(" SUBTYPES "" This is an additional protection mechanism over and above
96 that provided by the mandatory, privilege, and discretionary access
97 control systems. Each openable object may be associated with a
98 subtype, of which there are a fixed number at system generation time.
99 Any object of a non*null subtype may be accessed only by those processes
100 who have access rights to the subtype as well as the object. The
101 access right to a subtype is established by opening the subtype for
102 the desired access. Access to the subtype is granted or denied according
103 to the usual mandatory and discretionary rules. An object with
104 a non*null subtype can be accessed only when the open descriptor for
105 the subtype, to which access must already have been granted, is
106 presented, and when the desired access to the object is a subset of
107 the access granted to the subtype. This forms a mini*capability
108 mechanism with type extension, which is necessary for achieving access
109 control over objects such as directories. Thus there will be a
110 directory subtype, to prevent arbitrary programs from damaging the
111 directory system just because they have access to a particular
112 directory. The rules for subtype access occur in the open function and

```

```

113 all functions that require a subtype capability for accessing
114 an object.")
115
116
117 TYPES
118
119 $(FROM smx)
120 nonDisType: STRUCT_OF(
121     INTEGER securityLevel; SET_OF securityCat securityCatS;
122     INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
123 daType: SET_OF daMode;
124 modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
125 tiiStruct: STRUCT_OF(nonDisType nd; modeStruct ms; INTEGER owner, group;
126     SET_OF privType priv);
127
128 $(from fca " exportable)
129 openDescriptor: DESIGNATOR;
130 openModes: {omRead, omWrite, omExclusiveRead, omExclusiveWrite};
131 IOfunction: {rewind, etc}; $(names for special kinds of IO functions)
132 deviceType: {RK05, RWPO4, RWPO5, RWPO6, RSW04, TWE16, TM11, TU56, PR11,
133     PC11, LP11, IMPI1B, LHDH};
134 terminalGroup: DESIGNATOR;
135 fileNsp: {129 ..255}; $(nsp values for file systems)
136
137 $(from fca " redeclarable)
138 fileStatus: STRUCT_OF(INTEGER nBlocks, linkCount, timeLastMod;
139     seid subtype; BOOLEAN openAtCrash);
140 $(data about an openable object that is returned to the user)
141 globalData: STRUCT_OF(INTEGER linkCount, openCount, timeLastMod; seid subtype;
142     BOOLEAN openAtCrash);
143 $(state information for all openable objects)
144 fileBlock: VECTOR_OF CHAR;
145 ioStatus: STRUCT_OF(INTEGER devIndep, devDep);
146 $(result of an IO operation, including possible hardware failure)
147 openFileEntry: STRUCT_OF(seid openSeid; SET_OF openModes openMode);
148 $(entry in a process' open file table)
149 asyncId: CHAR;
150 readResult: STRUCT_OF(VECTOR_OF fileBlock data; ioStatus errst);
151 deviceStruct: STRUCT_OF(BOOLEAN addressable;
152     INTEGER minRequest, maxRequest, modSize, maxBlockNo);
153 $(properties necessary for processing IO requests for devices)
154 mountTableEntry: STRUCT_OF(seid devSeid; BOOLEAN readOnly;
155     tiiStruct devTii; globalData devGl);
156 fileSystemEntry: STRUCT_OF(seid fileSeid; globalData gl; tiiStruct tii;
157     VECTOR_OF fileBlock fileData);
158 $(the state of a file, for purposes of mounting and unmounting)
159 fileSystem: SET_OF fileSystemEntry;
160 $(the state of an entire mountable file system)
161 sodPair: STRUCT_OF(seid ps; openDescriptor od);
162 $(for openCount definition)
163
164
165 PARAMETERS
166
167 SET_OF seid subtypeSeidSet; $( set of non-null subtypes known to the system)
168 seid nullStSeid; $( seid indicating the null subtype)

```

```

169 openDescriptor FCAnullSt; $(a file descriptor indicating the null subtype)
170 seid FCARootSeid; $(distinguished root of KSOS permanent file system)
171 INTEGER FCAMaxOpenDescriptors; $(maximum number of open descriptors per
172     process)
173
174
175     DEFINITIONS
176
177 INTEGER FCAfileSize(seid fSeid) IS
178     CARDINALITY({INTEGER i | FCAfileData(fSeid, i) ~= ?});
179     $(the size, in blocks, of an addressable device, file, or extent)
180
181 INTEGER nOpenDescriptors(seid pSeid) IS
182     CARDINALITY({openDescriptor od | FCAopenEntry(pSeid, od) ~= ?});
183     $(the number of open objects in a given process; it must not exceed a fixed
184     maximum)
185
186 deviceStruct deviceDataSeid(seid fSeid) IS
187     DDFdeviceData(FCAdeviceType(fSeid));
188     $(given the seid of a device, the data on which its IO requests depend)
189
190 SET OF daMode h_modeTrans(SET OF openModes oModes) IS
191     (IF omRead INSET oModes THEN {daRead} ELSE {});
192     UNION (IF omWrite INSET oModes THEN {daWrite} ELSE {});
193     $(translates from one enumerated type, open Modes, to another slightly
194     different enumerated type, discretionary access modes)
195
196 BOOLEAN isCurrentPath(seid tSeid) IS
197     EXISTS terminalGroup t : FCACurrentPath(t) = tSeid;
198     $(for a given terminal, tells whether it is the active path to its physical
199     device)
200
201 BOOLEAN isReadOnly(seid s) IS
202     EXISTS fileNsp fsNsp
203         : SENseidNsp(FCAmountTable(fsNsp).devSeid) = SENseidNsp(s)
204           AND FCAmountTable(fsNsp).readOnly = TRUE;
205     $(tells whether or not a given file is on a file system that is mounted in
206     read-only mode)
207
208 $(When a file's link count goes to zero, its link count becomes '?', which
209 means that it does not exist for any operation, including K link. The
210 file cannot be physically deleted until its open count becomes zero, though.)
211
212 BOOLEAN logicalFileExistence(seid fSeid) IS
213     FCAinfo(fSeid) ~= ? AND FCAinfo(fSeid).linkCount ~= ?;
214
215     EXTERNALREFS
216
217     FROM mac:
218     VFUN MACclock() ^> INTEGER time;
219
220     FROM smx:
221     seid: DESIGNATOR;
222     secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
223         tExtent, tNull};
224     privType: {

```

```

225             privFileUpdateStatus, privLink,
226             privLockSeg,           privModifyPriv,
227             privMount,             privSetLevel,
228             privStickySeg,         privSetPath,
229             privViolSimpSecurity,  privViolStarSecurity,
230             privViolSimpIntegrity, privViolStarIntegrity,
231             privViolDiscrAccess,  privRealizeExecPermission,
232             privSignal,           privWalkPTable,
233             privHalt,             privKernelCall,
234             privViolCompartments, privSetComm,
235             privImmigrate
236         };
237
238 daMode: {daRead, daWrite, daExecute};
239 securityCat: DESIGNATOR;
240 integrityCat: DESIGNATOR;
241 VFUN SENSEidNsp(seid s) ^> INTEGER nsp;
242 VFUN SENSEidType(seid s) ^> secureEntityType set;
243 VFUN TIIinfo(seid s) ^> tiiStruct tlist;
244 VFUN SMXhasPriv(seid pSeid; privType priv) ^> BOOLEAN b;
245 VFUN SMXflow(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
246 VFUN SMXdap(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
247
248 FROM ddf:
249 VFUN DDFdeviceData(deviceType d) ^> deviceStruct ds;          $(DDFdeviceData)
250
251
252
253     ASSERTIONS
254
255 FORALL seid s; openDescriptor od
256     : SMXflow(s, FCAopenEntry(s, od).openSeid, {daRead});
257     $(FCAopenEntry only contains seids of objects whose security
258     levels are lower than that of the opening process)
259
260 FORALL seid pSeid;
261     openDescriptor od | FCAopenEntry(pSeid, od) ~= ?;
262     openFileEntry oe = FCAopenEntry(pSeid, od)
263     : (omWrite INSET oe.openMode
264     AND SENSEidType(oe.openSeid) ~= tSubtype)
265     => SMXflow(pSeid, oe.openSeid, {daRead, daWrite});
266     $( any object open for reading and writing that is not a subtype
267     must be at the same security level as the process that opened
268     it)
269 FORALL seid pSeid;
270     openDescriptor od | FCAopenEntry(pSeid, od) ~= ?;
271     openFileEntry oe = FCAopenEntry(pSeid, od)
272     : {omExclusiveRead, omExclusiveWrite} SUBSET oe.openMode
273     => SMXflow(pSeid, oe.openSeid, {daRead, daWrite});
274     $(all files open for exclusive use are open only by processes of
275     the same security level)
276
277 FORALL seid s | FCAinfo(s) ~= ? AND FCAinfo(s).subtype ~= nullStSeid
278     : SMXflow(s, FCAinfo(s).subtype, {daRead});
279     $( the level of a subtype is less than or equal to the levels
280     of any objects of that subtype)

```

```

281
282 FORALL seid fSeid | FCAinfo(fSeid) ~= ?
283 : SENSEidType(fSeid) INSET {tTerminal, tDevice, tSubtype, tFile, tExtent};
284 $(restricts the types of objects manipulated by this module)
285
286 FORALL seid fSeid: {INTEGER i | FCAfileData(fSeid, i) ~= ?}
287 = {0 .. FCAfileSize(fSeid) - 1};
288 $(the blocks of a file, extent, or addressable device form a sequence)
289
290 FORALL seid dSeid
291 | SENSEidType(dSeid) = tDevice AND FCAinfo(dSeid) ~= ?
292 : (LET deviceStruct d = deviceDataSeid(dSeid)
293 IN NOT d.addressable => d.modSize = 1 AND d.maxBlockNo = 0);
294 $(necessary properties of all non-addressable devices)
295
296 FORALL seid fSeid
297 | FCAinfo(fSeid) ~= ?
298 : (LET secureEntityType t = SENSEidType(fSeid)
299 IN FORALL INTEGER i | FCAfileData(fSeid, i) ~= ?
300 : LENGTH(FCAfileData(fSeid, i))
301 = (IF t = tDevice THEN deviceDataSeid(fSeid).modSize
302 ELSE IF t INSET {tExtent, tFile} THEN 512
303 ELSE ?));
304 $(the lengths of all blocks for files, extents, or addressable devices
305 must correspond to the parameters for those objects)
306
307 FORALL seid s | SENSEidType(s) = tTerminal AND FCAinfo(s) ~= ?
308 : (EXISTS terminalGroup t1
309 : s INSET FCAterminalPathSet(t1)
310 AND (FORALL terminalGroup t2 ~= t1
311 : NOT s INSET FCAterminalPathSet(t2)));
312 $(each terminal is in exactly one terminal group)
313
314 FORALL seid f
315 | EXISTS INTEGER i : FCAfileData(f, i) ~= ?
316 : FCAinfo(f) ~= ?
317 AND (SENSEidType(f) = tDevice AND deviceDataSeid(f).addressable = TRUE
318 OR SENSEidType(f) INSET {tFile, tExtent});
319 $(only files, extents, or addressable devices have file data)
320
321 FORALL seid f
322 : (FCAinputStream(f) ~= ? AND FCAoutputStream(f) ~= ?)
323 = (SENSEidType(f) = tTerminal
324 OR SENSEidType(f) = tDevice
325 AND deviceDataSeid(f).addressable = FALSE);
326 $(non-addressable device have both an input and an output stream, although
327 it may always be null)
328
329 FORALL seid pSeid | FCAopenTableExists(pSeid)
330 : FCAopenEntry(pSeid, FCAnullSt) = ?;
331 $(there are never any opened objects assigned to the open descriptor
332 reserved for the null subtype)
333
334 FORALL seid s1; seid s2
335 : SENSEidType(s1) = tSubtype AND SENSEidType(s2) = tSubtype
336 => SENSEidNsp(s1) = SENSEidNsp(s2);

```

```

337 $(subtypes have only one name space partition)
338
339 SENSEidType(nullStSeid) = tSubtype;
340 $(the null subtype seid is a subtype )
341
342 FORALL seid s INSET subtypeSeidSet : SENSEidType(s) = tSubtype;
343 $(properties of the set of non^null subtypes)
344
345 SENSEidType(FCARootSeid) = tFile;
346 $(property of the distinguished root of the entire file system)
347
348 FORALL seid f | FCAinfo(f) ~= ? AND SENSEidType(f) = tFile
349 : EXISTS fileNsp fsNsp : SENSEidNsp(FCAmountTable(fsNsp).devSeid)
350 = SENSEidNsp(f)
351 OR f = FCARootSeid;
352 $(a file is either the root or it is on a mountable file system)
353
354
355 FUNCTIONS
356
357 $(***** state functions ** devices *****)
358
359 VFUN FCADeviceType(seid fSeid) ^> deviceType d; $(FCADeviceType)
360 $(gives the type of a particular device, which determines its IO behavior)
361 HIDDEN;
362 INITIALLY
363 (d ~= ?)
364 = (SENSEidType(fSeid) = tDevice AND FCAinfo(fSeid) ~= ?);
365
366 $(***** state functions ** terminals *****)
367
368 VFUN FCATerminalPathSet(terminalGroup t) ^> SET_OF seid ss;
369 $(FCATerminalPathSet)
370 $(defines the set of paths, or logical terminals, that correspond to
371 a given terminal group, or physical terminal)
372 HIDDEN;
373 INITIALLY
374 FORALL seid s INSET ss
375 : FCAinfo(s) ~= ? AND SENSEidType(s) = tTerminal;
376
377 VFUN FCAcurrentPath(terminalGroup t) ^> seid s; $(FCAcurrentPath)
378 $(the seid of the logical terminal that is allowed to do IO on a given
379 physical terminal)
380 HIDDEN;
381 INITIALLY
382 s INSET FCATerminalPathSet(t);
383
384 $(***** state functions ** mountable file systems *****)
385
386 VFUN FCAMountTable(fileNsp fsNsp) ^> mountTableEntry mte; $(FCAMountTable)
387 $(for a given extent that is mounted, tells leaf of the old file system,
388 the root of the new or mounted file system, whether the file system
389 is read only, and the state information for the extent)
390 HIDDEN;
391 INITIALLY mte = ?;
392

```



```

393 VFUN FCAextentToFileSys(VECTOR_OF fileBlock fb; seid rootSeid)
394   ^> fileSystem fs;                                     $(FCAextentToFileSys)
395   $(given the data of a given extent and a root seid for a file system
396   to be made out of the extent whose data is given, produces the
397   file system, consisting of a set of tuples each made up of a
398   file seid and the state of the file)
399   HIDDEN;
400   INITIALLY
401     fs ~= ? =>
402       (EXISTS fileSystemEntry fse INSET fs : rootSeid = fse.fileSeid)
403       AND (FORALL fileSystemEntry fse INSET fs
404           : SENSEidNsp(fse.fileSeid) = SENSEidNsp(rootSeid));
405       $(a constant function for data conversion)
406
407   $(***** state functions ** files *****)
408
409 VFUN FCAinfo(seid fSeid) ^> globalData gl;               $(FCAinfo)
410   $(the status information, excluding data and type dependent stuff, for
411   an openable object)
412   HIDDEN;
413   INITIALLY
414     IF deviceDataSeid(fSeid) ~= ? THEN gl ~= ?
415     ELSE IF fSeid INSET {FCARootSeid, nullStSeid}
416       THEN gl.subtype = nullStSeid
417     ELSE IF fSeid INSET subtypeSeidSet THEN gl.subtype = nullStSeid
418     ELSE gl = ?;
419
420 VFUN FCAfileData(seid fSeid; INTEGER blockNo) ^> fileBlock fb; $(FCAfileData)
421   $(the data contained on a file, extent or an addressable device)
422   DEFINITIONS
423     secureEntityType type IS SENSEidType(fSeid);
424     deviceStruct d
425       IS IF type = tDevice THEN deviceDataSeid(fSeid)
426       ELSE IF type INSET {tFile, tExtent}
427         THEN STRUCT(TRUE, 512, 65534, 512, FCAfileSize(fSeid)*1)
428       ELSE STRUCT(FALSE, 1, 255, 1, 0);
429   HIDDEN;
430   INITIALLY
431     (IF FCAinfo(fSeid) = ? OR type INSET {tTerminal, tSubtype}
432     OR NOT blockNo INSET {0 .. d.maxBlockNo} OR NOT d.addressable
433     THEN fb = ?
434     ELSE fb ~= ?)
435     AND (fb ~= ?
436         => LENGTH(fb) = d.modSize
437         AND (FORALL INTEGER i INSET {0 .. blockNo - 1}
438             : FCAfileData(fSeid, i) ~= ?)
439         AND (FORALL INTEGER j INSET {1 .. LENGTH(fb)}
440             : fb[j] ~= ?));
441
442 VFUN FCAinputStream(seid fSeid) ^> VECTOR OF CHAR vc;    $(FCAinputStream)
443   $(the input data for terminals and nonaddressable devices)
444   HIDDEN;
445   INITIALLY
446     IF FCAinfo(fSeid) ~= ?
447       AND (SENSEidType(fSeid) = tTerminal
448       OR SENSEidType(fSeid) = tDevice

```

```

449             AND deviceDataSeid(fSeid).addressable = TRUE)
450     THEN vc ~= ?
451     ELSE vc = ?;
452
453 VFUN FCAoutputStream(seid fSeid) "> VECTOR_OF CHAR vc;    $(FCAoutputStream)
454 $(the output data for terminals and nonaddressable devices)
455     HIDDEN;
456     INITIALLY
457         IF FCAinfo(fSeid) ~= ?
458             AND (SENseidType(fSeid) = tTerminal
459                 OR SENseidType(fSeid) = tDevice
460                 AND deviceDataSeid(fSeid).addressable = TRUE)
461             THEN vc ~= ?
462             ELSE vc = ?;
463
464 $(***** state functions **** open tables *****)
465
466 VFUN FCAopenTableExists(seid pSeid) "> BOOLEAN b;    $(FCAopenTableExists)
467 $(the existence predicate for the table of openable objects corresponding
468     to a given process)
469     HIDDEN;
470     INITIALLY b = FALSE;
471
472 VFUN FCAopenEntry(seid pSeid; openDescriptor od) "> openFileEntry oe;
473 $(the information in entry "od" of the open table of process "pSeid")
474     HIDDEN;    $(FCAopenEntry)
475     INITIALLY oe = ?;
476
477 $(***** creation of files *****)
478
479 OVFUN FCAcreate(seid pSeid, nspSeid; SET_OF openModes mode; modeStruct d; openDescriptor st
480     "> STRUCT_OF(seid fSeid; openDescriptor od) return;    $(FCAcreate)
481
482 $(creates a file and opens it in the desired mode.  the file system on
483     which the file resides must be writable.  If a subtype capability
484     is provided, it must refer to a valid subtype.  The created file gets
485     only the discretionary access permission specified.  All other data
486     comes from the process making the call.  The file is originally of zero
487     length)
488     DEFINITIONS
489         tiiStruct ptii IS TIIinfo(pSeid);
490         tiiStruct otii IS STRUCT(ptii.nd, d, ptii.owner, ptii.group, ptii.priv);
491         seid extSeid
492             IS SOME seid s | SENseidNsp(FCAmountTable(SENseidNsp(s)).devSeid)
493                 = SENseidNsp(nspSeid);
494     EXCEPTIONS
495         $(these exceptions subsume those for opening a file for writing)
496         XbadModes: NOT ((mode = {omRead})
497             OR (mode = {omWrite})
498             OR (mode = {omRead, omWrite})
499             OR (mode = {omRead, omWrite, omExclusiveWrite})
500             OR (mode = {omRead, omWrite, omExclusiveRead, omExclusiveWrite}));
501         XbadStCap:
502             stCap ~= FCAnullSt AND SENseidType(FCAopenEntry(pSeid, stCap).openSeid) ~= tSubtype;
503
504         XodSpace: nOpenDescriptors(pSeid) >= FCAMaxOpenDescriptors;

```

```

505
506 XnoFile:                                $( actually means no file system )
507     NOT(extSeid ~= ? AND SMXflow(pSeid, extSeid, {daRead}));
508     $(the flow violation applies to the file system on which the newly
509     created file is to reside; its TII exists but not its global data)
510 XnotWritable: isReadOnly(nspSeid);
511 XbadSubtype:                             $( NOT currently checked in implementation JN)
512     stCap ~= FCAnullSt
513     AND NOT omWrite INSET FCAopenEntry(pSeid, stCap).openMode;
514 RESOURCE_ERROR;
515 ASSERTIONS
516     FCAopenTableExists(pSeid);
517     SMXflow(pSeid, nspSeid, {daRead});
518 EFFECTS
519     LET seid fs | FCAinfo(fs) = ?;
520     openDescriptor o | FCAopenEntry(pSeid, o) = ?
521     AND o ~= FCAnullSt
522     AND SENseidType(fs) = tFile
523     AND SENseidNsp (fs) = SENseidNsp(nspSeid)
524     IN 'TIIinfo(fs) = otii
525     AND 'FCAopenEntry(pSeid, o) = STRUCT(fs, mode)
526     AND return = STRUCT(fs,o)
527     AND 'FCAinfo(fs) = STRUCT(0, 1, MACclock(),
528                                IF stCap = FCAnullSt THEN nullStSeid
529                                ELSE FCAopenEntry(pSeid, stCap).openSeid,
530                                FALSE);
531
532 $(***** file opening and closing *****)
533
534 OVFUN FCAopen(seid pSeid, o; SET_OF openModes mode; openDescriptor stCap)
535     "> openDescriptor od;                                $(FCAopen)
536     $(opens the openable object specified by "o". The object must exist
537     and the process "pSeid" must have the right mandatory and discretionary
538     access. Special rules, described below, apply to subtypes.
539     Special rules also apply when the object is being opened for exclusive
540     use. Only one process may open an object for exclusive use.
541     A process is allowed a fixed maximum number of open objects.)
542 DEFINITIONS
543     globalData ofst IS FCAinfo(o);
544     openFileEntry stEntry IS FCAopenEntry(pSeid, stCap);
545
546 EXCEPTIONS
547     XbadModes: NOT ((mode = {omRead})                $( bad user argument )
548                    OR (mode = {omWrite})
549                    OR (mode = {omRead, omWrite})
550                    OR (mode = {omRead, omWrite, omExclusiveWrite})
551                    OR (mode = {omRead, omWrite, omExclusiveRead, omExclusiveWrite}));
552
553     XbadStCap:                                     $( bad subtype od given )
554     stCap ~= FCAnullSt AND SENseidType(stEntry.openSeid) ~= tSubtype;
555     $(a subtype capability was specified, but it is not for a valid subtype)
556
557     XodSpace: nOpenDescriptors(pSeid) >= FCAMaxOpenDescriptors;
558                                     $( no open descriptors left )
559
560     XnoFile: (ofst = ?)
561                                     $( inaccessible )

```

```

561      OR (omRead INSET mode AND NOT SMXflow(pSeid, o, {daRead }))
562      OR (omWrite INSET mode AND NOT SMXflow(pSeid, o, {daWrite}))
563      OR (SENseidType(o) = tFile AND NOT logicalFileExistence(o));
564                                     $( discretionary access ck )
565 XdapViol: NOT SMXdap(pSeid, o, h_modeTrans(mode));
566
567 XbadSubtype1:
568   stCap = FCAnullSt AND NOT ofst.subtype = nullStSeid;
569   $(no subtype capability was specified, but the object has a non^null
570     subtype)
571
572 XbadSubtype2:                                     $( unwanted subtype seid )
573   stCap ^= FCAnullSt AND ofst.subtype = nullStSeid;
574
575 XbadSubtype3:
576   stCap ^= FCAnullSt
577   AND (stEntry.openSeid ^= ofst.subtype   $( wrong subtype )
578       OR NOT (                                     $( open modes must be subset of subtype modes )
579           (mode INTER {omRead, omWrite})
580           SUBSET (stEntry.openMode INTER {omRead, omWrite})
581       )
582   );
583   $(a subtype capability is specified, but it does not match the subtype
584     of the object, with access modes included;
585     note that execute mode on the subtype is required to write on
586     an object of the subtype)
587 XnotWritable: omWrite INSET mode AND isReadOnly(o);
588 Xbusy:                                     $( exclusive use conflict )
589 EXISTS seid pSeidl; openDescriptor odl.
590 FCAopenEntry(pSeidl, odl).openSeid = o AND
591   (   (omRead INSET mode
592       AND omExclusiveRead INSET FCAopenEntry(pSeidl, odl).openMode)
593   OR (omWrite INSET mode
594       AND omExclusiveWrite INSET FCAopenEntry(pSeidl, odl).openMode)
595   OR (omExclusiveRead INSET mode
596       AND omRead INSET FCAopenEntry(pSeidl, odl).openMode)
597   OR (omExclusiveWrite INSET mode
598       AND omWrite INSET FCAopenEntry(pSeidl, odl).openMode)
599   );
600 ASSERTIONS
601   FCAopenTableExists(pSeid);
602 EFFECTS
603   LET openDescriptor odl | FCAopenEntry(pSeid, odl) = ? AND odl ^= FCAnullSt
604   IN 'FCAopenEntry(pSeid, odl) = STRUCT(o, mode)
605   AND od = odl;
606   'FCAinfo(o) = STRUCT(ofst.linkCount, ofst.openCount + 1,
607                       ofst.timeLastMod, ofst.subtype, ofst.openAtCrash);
608
609 OFUN FCAclose(seid pSeid; openDescriptor od);                                     $(FCAclose)
610   $(closes the open object named by "od". If the object is not is use by
611     anyone else, either by linking or by opening, then the object is deleted)
612 DEFINITIONS
613   openFileEntry oe IS FCAopenEntry(pSeid, od);
614   seid fSeid IS oe.openSeid;
615   globalData ofst IS FCAinfo(fSeid);
616 EXCEPTIONS

```

```

617     XbadOd: oe = ?;
618     ASSERTIONS
619     FCAopenTableExists(pSeid);
620     EFFECTS
621     IF (ofst.linkCount = 0 OR ofst.linkCount = ?)
622         AND ofst.openCount = 1
623         AND SENSEIDType(fSeid) = tFile
624     THEN 'FCAinfo(fSeid) = ?
625         AND (FORALL INTEGER i : 'FCAfileData(fSeid, i) = ?)
626         AND 'TIIinfo(fSeid) = ?
627     ELSE 'FCAinfo(fSeid) = STRUCT(ofst.linkCount, ofst.openCount + 1,
628                                     ofst.timeLastMod, ofst.subtype,
629                                     ofst.openAtCrash);
630     'FCAopenEntry(pSeid, od) = ?;
631     $(the openAtCrash field is cleared when closing an object that was open
632     for writing. This has not been put in.)
633
634     $(***** open table maintenance *****)
635
636     OFUN FCAcreateOpenTable(seid pSeid);                                $(FCAcreateOpenTable)
637     $( this operation creates an open table associated with a process seid;
638     it is used as an auxiliary operation by the process modules when
639     creating a new a process)
640     ASSERTIONS
641     NOT FCAopenTableExists(pSeid);
642     EFFECTS
643     'FCAopenTableExists(pSeid) = TRUE;
644
645     OFUN FCAdeleteOpenTable(seid pSeid);                                $(FCAdeleteOpenTable)
646     $( Deletes the open table associated with a process; supports the release
647     of a process)
648     ASSERTIONS
649     FCAopenTableExists(pSeid);
650     EFFECTS
651     'FCAopenTableExists(pSeid) = FALSE;
652
653     $(***** utility operations *****)
654
655     OFUN FCAcopyOpenTable(seid fromSeid, toSeid);                        $(FCAcopyOpenTable)
656     $(Copies the contents of one open table, "fromSeid," to another, "toSeid."
657     "toSeid" must be empty)
658     ASSERTIONS
659     FCAopenTableExists(fromSeid);
660     FCAopenTableExists(toSeid);
661     FORALL openDescriptor od : FCAopenEntry(toSeid, od) = ?;
662     EFFECTS
663     FORALL openDescriptor od;
664         seid fSeid = FCAopenEntry(fromSeid, od).openSeid;
665         globalData ofst = FCAinfo(fSeid)
666     : 'FCAopenEntry(toSeid, od) = FCAopenEntry(fromSeid, od)
667         AND (ofst ~= ?
668             => 'FCAinfo(fSeid)
669                 = STRUCT(ofst.linkCount, ofst.openCount + 1,
670                         ofst.timeLastMod, ofst.subtype,
671                         ofst.openAtCrash));
672

```

```

673 $(***** file status operations *****)
674
675 VFUN FCAGetFileStatus(seid pSeid, fSeid) #> fileStatus fst;$(FCAGetFileStatus)
676 $(returns the status of the file. The requesting process must have
677   mandatory access to the object, and the object must exist)
678 DEFINITIONS
679   seid f IS fSeid;
680   globalData gl IS FCAinfo(f);
681 EXCEPTIONS
682   XnoFile: NOT(FCAinfo(f) ~= ? AND SMXflow(pSeid, f, {daRead}));
683 DERIVATION
684   STRUCT(IF FCAfileSize(f) = ? THEN 0 ELSE FCAfileSize(f),
685     gl.linkCount, gl.timeLastMod, gl.subtype,
686     gl.openAtCrash);
687
688 END_MODULE


```

AD-A111 578

FORD AEROSPACE AND COMMUNICATIONS CORP PALO ALTO CA W--ETC F/8 9/2
KSOS COMPUTER PROGRAM DEVELOPMENT SPECIFICATIONS (TYPE B-5), (K--ETC(U)
DEC 80 MDA903-77-C-0333
MDL-TR7932-REV-3 NL

UNCLASSIFIED


2 of 4
20
01/10/76



SSIF 120

2 OF 4

AD
A1111578




```

1  $("  MODULE:      fmi.specs (version 4.14)
2      CONTENTS:    File Miscellaneous Operations
3      TYPE:        SPECIAL specifications
4      LAST CHANGED: 12/23/80, 16:48:40
5  ")
6
7
8  MODULE fmi
9
10
11  $( this module contains miscellaneous file operations that are not included
12    in the fca module, because the SPECIAL checker at FACC cannot accomodate
13    the combined file)
14
15    TYPES
16
17    $(FROM smx)
18    nonDisType: STRUCT_OF(
19        INTEGER securityLevel; SET_OF securityCat securityCatS;
20        INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
21    daType: SET_OF daMode;
22    modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
23    tiiStruct: STRUCT_OF(nonDisType nd; modeStruct ms; INTEGER owner, group;
24        SET_OF privType priv);
25
26    $(from fca)
27    globalData: STRUCT_OF(INTEGER linkCount, openCount, timeLastMod;
28        seid subtype; BOOLEAN openAtCrash);
29    ioStatus: STRUCT_OF(INTEGER devIndep, devDep);
30    asyncId: CHAR;
31    fileBlock: VECTOR_OF CHAR;
32    openFileEntry: STRUCT_OF(seid openSeid; SET_OF openModes openMode);
33    readResult: STRUCT_OF(VECTOR_OF fileBlock data; ioStatus errst);
34    deviceStruct: STRUCT_OF(BOOLEAN addressable;
35        INTEGER minRequest, maxRequest, modSize, maxBlockNo);
36    mountTableEntry: STRUCT_OF(seid devSeid; BOOLEAN readOnly;
37        tiiStruct devTii; globalData devGl);
38    fileSystemEntry: STRUCT_OF(seid fileSeid; globalData gl; tiiStruct tii;
39        VECTOR_OF fileBlock fileData);
40    fileSystem: SET_OF fileSystemEntry;
41
42    fileNsp: {129 .. 255};          $( nsp values for file systems )
43
44    DEFINITIONS
45
46    $(these definitions are explained in the fca module)
47
48    INTEGER FCAfileSize(seid fSeid) IS
49    CARDINALITY({INTEGER i | FCAfileData(fSeid, i) ~= ?});
50
51    deviceStruct deviceDataSeid(seid fSeid) IS
52    DDFdeviceData(FCAdeviceType(fSeid));
53
54    BOOLEAN isCurrentPath(seid tSeid) IS
55    EXISTS terminalGroup t : FCAcurrentPath(t) = tSeid;
56

```

```

57 BOOLEAN isReadOnly(seid fSeid) IS
58     FCAmountTable(SENseidNsp(fSeid)).readOnly = TRUE;
59
60 $(When a file's link count goes to zero, its link count becomes '?', which
61 means that it does not exist for any operation, including K_link. The
62 file cannot be physically deleted until its open count becomes zero, though.)
63
64 BOOLEAN logicalFileExistence(seid fSeid) IS
65     FCAinfo(fSeid) ~= ? AND FCAinfo(fSeid).linkCount ~= ?;
66
67     EXTERNALREFS
68
69     FROM ddf:
70
71 VFUN DDFdeviceData(deviceType d) ^> deviceStruct ds;
72
73     FROM smx:
74 seid: DESIGNATOR;
75 secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
76                    tExtent, tNull};
77     privType: {
78                 privFileUpdateStatus, privLink,
79                 privLockSeg,          privModifyPriv,
80                 privMount,            privSetLevel,
81                 privStickySeg,        privSetPath,
82                 privViolSimpSecurity, privViolStarSecurity,
83                 privViolSimpIntegrity, privViolStarIntegrity,
84                 privViolDiscrAccess,  privRealizeExecPermission,
85                 privSignal,           privWalkPTable,
86                 privHalt,              privKernelCall,
87                 privViolCompartments, privSetComm,
88                 privImmigrate
89     };
90 securityCat: DESIGNATOR;
91 integrityCat: DESIGNATOR;
92 daMode: {daRead, daWrite, daExecute};
93 VFUN SENseidNsp(seid s) ^> INTEGER nsp;
94 VFUN SENseidType(seid s) ^> secureEntityType set;
95 VFUN TIIinfo(seid s) ^> tiiStruct tiist;
96 VFUN SMXhasPriv(seid pSeid; privType priv) ^> BOOLEAN b;
97 VFUN SMXflow(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
98 VFUN SMXdap(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
99
100    FROM fca:
101 openDescriptor: DESIGNATOR;
102 openModes: {omRead, omWrite, omExclusiveRead, omExclusiveWrite};
103 IOfunction: {rewind, etc};
104 deviceType: {RK05, RWPO4, RWPO5, RWPO6, RSWO4, TWE16, TM11, TU56, PR11,
105             PC11, LP11, IMPL1B, LHDH};
106 terminalGroup: DESIGNATOR;
107 seid FCARootSeid;
108 VFUN FCAdeviceType(seid fSeid) ^> deviceType d;
109 VFUN FCAcurrentPath(terminalGroup t) ^> seid s;
110 VFUN FCAterminalPathSet(terminalGroup t) ^> SET_OF seid ss;
111 VFUN FCAmountTable(fileNsp nsp) ^> mountTableEntry mte;
112 VFUN FCAextentToFileSys(VECTOR_OF fileBlock fb; fileNsp nsp)

```

```

113  *> fileSystem fs;
114 VFUN FCAinfo(seid fSeid) *> globalData gl;
115 VFUN FCAfileData(seid fSeid; INTEGER blockNo) *> fileBlock fb;
116 VFUN FCAinputStream(seid fSeid) *> VECTOR OF CHAR vc;
117 VFUN FCAoutputStream(seid fSeid) *> VECTOR OF CHAR vc;
118 VFUN FCAopenTableExists(seid pSeid) *> BOOLEAN b;
119 VFUN FCAopenEntry(seid pSeid; openDescriptor od) *> openFileEntry oe;
120
121
122  FUNCTIONS
123
124  $(***** process support operations *****)
125
126  $(Note: the function FCAcreateOpenTable is sufficient to support PROspawn.
127  The other support operations are listed here.)
128
129 OFUN FMiforkSupport(seid parent, child);                                $(FMiforkSupport)
130  $(This function creates a new open table, "child," and copies all open
131  from "parent" into it)
132  EXCEPTIONS
133    XexclusiveFile:
134      EXISTS openDescriptor od
135        : {omExclusiveRead,omExclusiveWrite} SUBSET
136          FCAopenEntry(parent, od).openMode;
137  ASSERTIONS
138    FCAopenTableExists(parent);
139    NOT FCAopenTableExists(child);
140  EFFECTS
141    'FCAopenTableExists(child) = TRUE;
142    FORALL openDescriptor od;
143      openFileEntry oe = FCAopenEntry(parent, od);
144      globalData gl = FCAinfo(oe.openSeid)
145      : 'FCAopenEntry(child, od) = oe
146        AND 'FCAinfo(oe.openSeid)
147          = STRUCT(gl.linkCount, gl.openCount + 1, gl.timeLastMod,
148          gl.subtype, gl.openAtCrash);
149
150
151 OFUN FMireleaseSupport(seid pSeid);                                $(FCAreleaseSupport)
152  $( Closes all the open objects of an open object table and deletes the
153  table)
154  DEFINITIONS
155    INTEGER nOpen(openDescriptor od)
156    IS CARDINALITY({openDescriptor odl
157      | FCAopenEntry(pSeid, od).openSeid
158      = FCAopenEntry(pSeid, odl).openSeid});
159    $(the number of times an object corresponding to a
160    given open descriptor is opened by this process)
161    seid s(openDescriptor od) IS FCAopenEntry(pSeid, od).openSeid;
162  ASSERTIONS
163    FCAopenTableExists(pSeid);
164  EFFECTS
165    FORALL openDescriptor od | FCAopenEntry(pSeid, od) ~= ?;
166      globalData gl = FCAinfo(s(od))
167      : 'FCAopenEntry(pSeid, od) = ?
168      AND (IF( gl.linkCount = 0 OR gl.linkCount =?))

```

```

169         AND gl.openCount = nOpen(od)
170         AND SENSeidType(s(od)) = tFile
171         THEN 'FCAinfo(s(od)) = ?
172             AND 'TIIinfo(s(od)) = ?
173         AND (FORALL INTEGER i : 'FCAfileData(s(od), i) = ?)
174         ELSE
175             'FCAinfo(s(od))
176             = STRUCT(gl.linkCount,
177                     gl.openCount * nOpen(od),
178                     gl.timeLastMod, gl.subtype,
179                     gl.openAtCrash));
180     'FCAopenTableExists(pSeid) = FALSE;
181
182 $(***** link and unlink operations *****)
183
184 OFUN FCAlink(seid pSeid, f);                                $(FCAlink)
185 $(increments the link count of an existing file. Mandatory access is
186   required, and the process must have the privilege to link. The file
187   system must not be mounted in read only mode)
188 DEFINITIONS
189   globalData fs IS FCAinfo(f);
190 EXCEPTIONS
191   XbadPriv: NOT SMXhasPriv(pSeid, privLink);
192   XnoFile: NOT(logicalFileExistence(f) AND SMXflow(pSeid, f, {daWrite}));
193   XnoFile: SENSeidType(f) ~= tFile;   $( not a file, but exists )
194   XnotWritable: isReadOnly(f);
195 ASSERTIONS
196   FCAopenTableExists(pSeid);
197 EFFECTS
198   'FCAinfo(f)
199   = STRUCT(fs.linkCount + 1, fs.openCount, fs.timeLastMod,
200           fs.subtype, fs.openAtCrash);
201
202 OFUN FCAunlink(seid pSeid, f);                                $(FCAunlink)
203 $(decrements the reference count of an existing file. Mandatory access is
204   required and the process must have the privilege to link. The file
205   system that the file is on must not be mounted in read only mode. Note
206   that unlinking can cause the file to be deleted if the link count goes
207   to zero and the file is not open anywhere)
208 DEFINITIONS
209   globalData fs IS FCAinfo(f);
210 EXCEPTIONS
211   XbadPriv: NOT SMXhasPriv(pSeid, privLink);
212   XnoFile: NOT(logicalFileExistence(f) AND SMXflow(pSeid, f, {daWrite}));
213   XnoFile: SENSeidType(f) ~= tFile;
214   XnotWritable: isReadOnly(f);
215 ASSERTIONS
216   FCAopenTableExists(pSeid);
217 EFFECTS
218   IF ((fs.linkCount <= 1) OR (fs.linkCount = ?))
219     AND fs.openCount = 0
220                                     $(if true, then delete file)
221   THEN 'FCAinfo(f) = ?
222       AND (FORALL INTEGER i : 'FCAfileData(f, i) = ?)
223       AND 'TIIinfo(f) = ?
224   ELSE 'FCAinfo(f) =

```

```

225      STRUCT((IF fs.linkCount <= 1 THEN ? ELSE fs.linkCount * 1),
226              fs.openCount, fs.timeLastMod, fs.subtype, fs.openAtCrash);
227 $(***** basic I/O operations *****)
228
229 VFUN FCAvReadBlocks(seid pSeid; openDescriptor od; INTEGER blockNo, size;
230                   asyncId id) $(FCAvReadBlocks)
231     > readResult rr;
232 $(the purpose of this function is to return the result that FCAreadBlocks
233    would return if executed)
234 $(returns blocks that are read from a given file, device, extent, or
235    terminal. the object must be open for reading. the block number and
236    size must be within range for the kind of object specified. note
237    that files, extents, and addressable devices are handled differently
238    from terminals and nonaddressable devices, with respect to the data.)
239 DEFINITIONS
240     seid fSeid IS FCAopenEntry(pSeid, od).openSeid;
241     deviceStruct d
242       IS IF SENSeidType(fSeid) = tDevice THEN deviceDataSeid(fSeid)
243         ELSE IF SENSeidType(fSeid) INSET {tFile, tExtent}
244           THEN STRUCT(TRUE, 512, 32768, 512, FCAfileSize(fSeid)*1)
245           ELSE STRUCT(FALSE, 1, 255, 1, 0); $(tTerminal)
246     INTEGER blockCount IS size/512;
247     INTEGER blockPast IS (IF SENSeidType(fSeid) = tDevice THEN
248                          MIN({d.maxBlockNo, blockNo+blockCount})
249                          ELSE
250                          MIN({FCAfileSize(fSeid), blockNo+blockCount}));
251
252 EXCEPTIONS
253     XbadOd: FCAopenEntry(pSeid, od) = ?;
254     XbadSize: size = 0;
255     XnotReadable: NOT omRead INSET FCAopenEntry(pSeid, od).openMode;
256     $( EXCEPTIONS AFTER THIS POINT ARE ACTUALLY HANDLED AS I/O ERRORS;
257       In addition, the following exceptions may not be treated as "pure"
258       in the implementation, i.e., they may involve effects in the code.)
259     XbadBlockNo:
260       NOT size INSET {d.minRequest .. d.maxRequest}
261       OR (size MOD d.modSize) /= 0;
262     XbadBlockNo: NOT blockNo INSET {0 .. d.maxBlockNo};
263     XendOfFile:
264       d.addressable AND blockNo + size/d.modSize > d.maxBlockNo;
265     Xerror: RESOURCE_ERROR;
266 ASSERTIONS
267     FCAopenTableExists(pSeid);
268 DERIVATION
269     IF d.addressable THEN
270       STRUCT(
271         VECTOR(FOR i FROM blockNo TO blockPast * 1
272                : FCAfileData(fSeid, i)),
273         SOME ioStatus ios | TRUE)
274     ELSE
275       LET INTEGER s1 INSET {0 .. size/d.modSize}
276       $(Must let s1 be nondeterministically selected for nonaddress. dev)
277       IN STRUCT(
278         VECTOR(FOR i FROM 1 TO s1 $(d.modSize = 1)
279                : VECTOR(FCAinputStream(fSeid)[i])),
280         SOME ioStatus ios | TRUE);

```

```

281
282 OVFUN FCAREadBlocks(seid pSeid; openDescriptor od; INTEGER blockNo, size;
283                      asyncId id)
284     "> readResult rr;                                     $(FCAREadBlocks)
285     $(LR ** needs semantics for asynchronous I/O)
286     $(returns blocks that are read from a given file, device, extent, or
287     terminal. the object must be open for reading. the block number and
288     size must be within range for the kind of object specified. note
289     that files, extents, and addressable devices are handled differently
290     from terminals and nonaddressable devices, with respect to the data.)
291     DEFINITIONS
292     seid fSeid IS FCAOpenEntry(pSeid, od).openSeid;
293     deviceStruct d
294         IS IF SENSEIDType(fSeid) = tDevice THEN deviceDataSeid(fSeid)
295         ELSE IF SENSEIDType(fSeid) INSET {tFile, tExtent}
296             THEN STRUCT(TRUE, 512, 32768, 512, FCAfileSize(fSeid)*1)
297             ELSE STRUCT(FALSE, 1, 255, 1, 0); $(tTerminal)
298     readResult rrl IS FCAvReadBlocks(pSeid, od, blockNo, size, id);
299     EXCEPTIONS
300     EXCEPTIONS_OF FCAvReadBlocks(pSeid, od, blockNo, size, id);
301     ASSERTIONS
302     FCAOpenTableExists(pSeid);
303     EFFECTS
304     rr = rrl;
305     NOT d.addressable
306     => 'FCAInputStream(fSeid)
307         = VECTOR(FOR i FROM LENGTH(rrl.data) + 1
308                 TO LENGTH(FCAInputStream(fSeid))
309                 : FCAInputStream(fSeid)[i]);
310
311 OVFUN FCAwriteBlocks(seid pSeid; openDescriptor od; INTEGER blockNo;
312                     VECTOR_OF fileBlock vfb; asyncId id)
313     "> ioStatus ios;                                     $(FCAwriteBlocks)
314     $(LR ** needs asynchronous I/O)
315     $(writes the contents of a vector of file blocks onto the object mentioned.
316     the data must correspond to the parameters of the openable object.)
317     DEFINITIONS
318     seid fSeid IS FCAOpenEntry(pSeid, od).openSeid;
319     secureEntityType type IS SENSEIDType(fSeid);
320     deviceStruct d
321         IS IF type = tDevice THEN deviceDataSeid(fSeid)
322         ELSE IF type INSET {tFile, tExtent}
323             THEN STRUCT(TRUE, 512, 32768, 512, FCAfileSize(fSeid))
324             ELSE STRUCT(FALSE, 1, 255, 1, 0);
325     INTEGER size IS LENGTH(vfb);
326     INTEGER blockCount IS size/512;
327     INTEGER blockPast IS (IF SENSEIDType(fSeid) = tDevice THEN
328                           MIN({d.maxBlockNo, blockNo+blockCount})
329                           ELSE
330                           blockNo+blockCount);
331
332     EXCEPTIONS
333     XbadOd: NOT (FCAOpenEntry(pSeid, od) ~= ?);
334     XbadSize: LENGTH(vfb) = 0;
335     XnotWritable: NOT omWrite INSET FCAOpenEntry(pSeid, od).openMode;
336     XnoSpace: RESOURCE_ERROR;

```

```

337 $( EXCEPTIONS AFTER THIS POINT ARE ACTUALLY HANDLED AS I/O ERRORS
338     In addition, the following exceptions may not be treated as "pure"
339     in the implementation, i.e., they may involve effects in the code.)
340     XbadSize:
341         NOT LENGTH(vfb) INSET {d.minRequest .. d.maxRequest}
342         OR (LENGTH(vfb) MOD d.modSize) ~= 0;
343     XbadBlockNo: NOT {blockNo .. blockNo + LENGTH(vfb) # 1} SUBSET
344         {0 .. d.maxBlockNo};
345     XendOfFile: d.addressable AND type ~= tFile
346         AND blockNo + LENGTH(vfb) > d.maxBlockNo;
347     XnoSpace: RESOURCE_ERROR;
348     ASSERTIONS
349     FORALL INTEGER i | vfb[i] ~= ?; INTEGER j | vfb[j] ~= ? :
350         LENGTH(vfb[i]) = LENGTH(vfb[j]);
351
352     EFFECTS
353     ios = (SOME ioStatus ios1 | TRUE);
354     IF d.addressable
355     THEN FORALL INTEGER i
356         : 'FCAfileData(fSeid, i)
357         = (IF NOT i INSET {blockNo .. blockPast # 1}
358             THEN FCAfileData(fSeid, i)
359             ELSE vfb[i # blockNo + 1])
360     ELSE 'FCAoutputStream(fSeid)
361         = VECTOR(FOR i FROM 1
362                 TO LENGTH(FCAoutputStream(fSeid))+LENGTH(vfb)
363                 : IF i <= LENGTH(FCAoutputStream(fSeid))
364                     THEN FCAoutputStream(fSeid)[i]
365                     ELSE
366                         vfb[(i#LENGTH(FCAoutputStream(fSeid)))/LENGTH(vfb)]
367                         [(i#LENGTH(FCAoutputStream(fSeid)))MOD LENGTH(vfb)]);
368
369
370 $(***** general device manipulation *****)
371
372 VFUN FCAvDeviceFunction(seid pSeid; openDescriptor od; IOfunction f;
373     VECTOR OF INTEGER args; asyncId id)
374     => ioStatus status; $(FCAvDeviceFunction)
375     $(the purpose of this function is to return the value that FCAdeviceFunction
376     would return if it were executed in a given state)
377     $(the specification of this function is device#dependent, but may be
378     filled in at a later time)
379     DEFINITIONS
380     seid dSeid IS FCAopenEntry(pSeid, od).openSeid;
381     EXCEPTIONS
382     XnoFile: FCAopenEntry(pSeid, od) = ?;
383     XbadDevice: NOT SENSEidType(dSeid) INSET {tDevice, tTerminal};
384     $( there should be an exception here which uses a table of functions
385     vs read/write permissions required to return XnotReadable or
386     XnotWritable as required. JN )
387     ASSERTIONS
388     FCAopenTableExists(pSeid);
389     DERIVATION
390     SOME ioStatus ios | TRUE;
391
392 OVFUN FCAdeviceFunction(seid pSeid; openDescriptor od; IOfunction f;

```

```

393                                VECTOR_OF INTEGER args; asyncId id)
394    ^> ioStatus status;                                $(FCAdeviceFunction)
395    $(the specification of this function is device^dependent, but may be
396    filled in at a later time)
397    EXCEPTIONS
398    EXCEPTIONS_OF FCAvDeviceFunction(pSeid, od, f, args, id);
399    ASSERTIONS
400    FCAopenTableExists(pSeid);
401    EFFECTS
402    $(the state of the device somehow changes, and the result is returned via
403    the io status)
404    status = FCAvDeviceFunction(pSeid, od, f, args, id);
405    TRUE;
406
407    OFUN FCAterminalLock(seid pSeid, devSeid);                $(FCAterminalLock)
408    $(sets the current terminal in the group to be "devSeid". The requesting
409    process must have the privilege to lock terminals)
410    $( Actually, this functionality is now performed by the
411    function SETPATH through K_device_function. JN)
412    EXCEPTIONS
413    XnoPriv: NOT SMXhasPriv(pSeid, privSetPath);
414    XnoClass: SENseidType(devSeid) ~= tTerminal;
415    XnoFile: FCAinfo(devSeid) = ?;
416    ASSERTIONS
417    FCAopenTableExists(pSeid);
418    EFFECTS
419    LET terminalGroup t | devSeid INSET FCAterminalPathSet(t)
420    IN 'FCAcurrentPath(t) = devSeid;
421
422    $(***** mounting and unmounting operations *****)
423
424    OVFUN FCAmount(seid pSeid, dev; BOOLEAN readOnly)    $(FCAmount)
425    ^> fileNsp newNsp;
426    $(performs the logical mounting of a file system from an extent. The
427    semantics are described in the general commentary in the FCA
428    specification)
429    DEFINITIONS
430    fileNsp openSlot IS
431    SOME fileNsp freeMountNsp | FCAmountTable(freeMountNsp) = ?;
432    fileSystem fileSys IS
433    FCAextentToFileSys(VECTOR(FOR i FROM 1 TO FCAfileSize(dev)
434    : FCAfileData(dev, i)), openSlot) ;
435    $(the file system produced by the data on the extent)
436    fileSystemEntry fse(seid f) IS
437    SOME fileSystemEntry fsel | fsel INSET fileSys AND fsel.fileSeid = f;
438    $(the entry in the file system with a given seid)
439    EXCEPTIONS
440    XbadPriv: NOT SMXhasPriv(pSeid, privMount);
441    XnoFile: NOT (FCAinfo(dev) ~= ?);
442    XnoClass: SENseidType(dev) ~= tExtent;
443    XmntSpace: RESOURCE_ERROR;
444    ASSERTIONS
445    FCAopenTableExists(pSeid);
446    EFFECTS
447    newNsp = openSlot;                                $( return new mount slot )
448    'FCAmountTable(openSlot)

```



```

449   = STRUCT(dev, readOnly, TIIinfo(dev), FCAinfo(dev));
450   'FCAinfo(dev) = ?;
451   FORALL INTEGER i : 'FCAfileData(dev, i) = ?;
452   'TIIinfo(dev) = ?;
453   FORALL seid f | fse(f) ~= ?
454   : 'FCAinfo(f) = fse(f).gl
455     AND (FORALL INTEGER i INSET {1 .. LENGTH(fse(f).fileData)}
456         : 'FCAfileData(f, i * 1) = fse(f).fileData[i])
457     AND 'TIIinfo(f) = fse(f).tii;
458
459   OFUN FCAunmount(seid pSeid; fileNsp fsNsp);                                $(FCAunmount)
460   $(performs logical unmounting of a file system. Restores the extent
461   so that it can be accessed)
462   DEFINITIONS
463   mountTableEntry mte IS FCAmountTable(fsNsp);
464   fileSystem fs IS
465     {fileSystemEntry fse
466     | LET seid fSeid | SENseidNsp(fSeid) = fsNsp AND FCAinfo(fSeid) ~= ?
467       IN fse = STRUCT(fSeid, FCAinfo(fSeid), TIIinfo(fSeid),
468         VECTOR(FOR i FROM 1 TO FCAfileSize(fSeid)
469             : FCAfileData(fSeid, i * 1)))};
470   $( the state of the file system to be unmounted)
471   VECTOR_OF fileBlock extData
472   IS SOME VECTOR_OF fileBlock vfb
473   | FCAextentToFileSys(vfb, fsNsp) = fs;
474   $(given the state of the file system, the extent that is equivalent to
475   it)
476   seid dev IS FCAmountTable(fsNsp).devSeid; $( underlying device )
477   EXCEPTIONS
478   XbadPriv: NOT SMXhasPriv(pSeid, privMount);
479   XnoXXX: FCAmountTable(fsNsp) = ?;
480   XnoTranquil:
481     EXISTS seid fSeid | SENseidNsp(fSeid) = fsNsp
482       : FCAinfo(fSeid).openCount > 0;
483   ASSERTIONS
484   FCAopenTableExists(pSeid);
485   EFFECTS
486   'FCAinfo(dev) = mte.devGl;
487   'TIIinfo(dev) = mte.devTii;
488   FORALL INTEGER i INSET {1 .. LENGTH(extData)}
489   : 'FCAfileData(dev, i * 1) = extData[i];
490   $(the device's state "comes back")
491   FORALL seid fSeid | SENseidNsp(fSeid) = fsNsp
492   : 'TIIinfo(fSeid) = ?
493     AND 'FCAinfo(fSeid) = ?
494     AND (FORALL INTEGER i : 'FCAfileData(fSeid, i) = ?);
495   $(the state of all files on the file system "disappears")
496   FCAmountTable(fsNsp) = ?;
497
498
499   END_MODULE

```

```

1  $( "      MODULE:      lev.specs (version 4.9)
2          CONTENTS:      System Levels
3          TYPE:          SPECIAL specifications
4          LAST CHANGED:   7/28/80, 09:51:43
5      ")
6
7
8  MODULE lev
9
10
11  $( This module enables the centralization of the get and set level operations
12    for all modules. Each module maintains the typeaindependent operation
13    of its own objects, and applies certain conditions to the getting and
14    setting of this information. However, there is only one kernel operation
15    for getting levels, and one for setting levels, of all objects. The
16    operations are combined in this module)
17
18    TYPES
19
20    $(from smx)
21  nonDisType: STRUCT_OF(
22              INTEGER securityLevel; SET_OF securityCat securityCatS;
23              INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
24  daType: SET_OF daMode;
25  modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
26  tiiStruct: STRUCT_OF(nonDisType nd; modeStruct ms; INTEGER owner, group;
27                      SET_OF privType priv);
28
29    $(FROM pvm)
30  globalData:
31    STRUCT_OF(BOOLEAN sharable, swappable, sticky, memAdvise, executable;
32              direction growth);
33  instanceStruct:
34    STRUCT_OF(globalData gl; direction growth; INTEGER refCount; VECTOR_OF INTEGER data);
35
36    $(FROM fca)
37  openFileEntry: STRUCT_OF(seid openSeid; SET_OF openModes openMode);
38
39
40    EXTERNALREFS
41
42    FROM smx:
43  seid: DESIGNATOR;
44  secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
45                    tExtent, tNull};
46  securityCat: DESIGNATOR;
47  integrityCat: DESIGNATOR;
48  daMode: {daRead, daWrite, daExecute};
49  privType: {
50
51              privFileUpdateStatus,  privLink,
52              privLockSeg,           privModifyPriv,
53              privMount,             privSetLevel,
54              privStickySeg,         privSetPath,
55              privViolSimpSecurity,  privViolStarSecurity,
56              privViolSimpIntegrity, privViolStarIntegrity,
57              privViolDiscrAccess,   privRealizeExecPermission,

```

```

57             privSignal,             privWalkPTable,
58             privHalt,               privKernelCall,
59             privViolCompartments,   privSetComm,
60             privImmigrate
61         };
62
63 VFUN SENseidType(seid s) ^> secureEntityType set;
64 VFUN TIIgetEntityLevel(seid pSeid, oSeid) ^> tiiStruct ntii;
65 OFUN TIIsetEntityLevel(seid pSeid, oSeid; tiiStruct ntii);
66
67     FROM pvm:
68 direction: {up, down};
69 VFUN SEGinstanceInfo(seid s) ^> instanceStruct is;
70
71     FROM fca:
72 openDescriptor: DESIGNATOR;
73 openModes: {omRead, omWrite, omExclusiveRead, omExclusiveWrite};
74 VFUN FCAopenEntry(seid pSeid; openDescriptor od) ^> openFileEntry oe;
75
76
77     FUNCTIONS
78
79 VFUN LEVgetObjectLevel(seid pSeid, oSeid) ^> tiiStruct otii;
80 EXCEPTIONS
81     EXCEPTIONS_OF TIIgetEntityLevel(pSeid, oSeid);
82 DERIVATION
83     TIIgetEntityLevel(pSeid, oSeid);
84
85 OFUN LEVsetObjectLevel(seid pSeid, oSeid; tiiStruct otii);
86                                     $(LEVsetObjectLevel)
87
88     DEFINITIONS
89     secureEntityType type IS SENseidType(oSeid);
90 EXCEPTIONS
91     EXCEPTIONS_OF TIIsetEntityLevel(pSeid, oSeid, otii);
92     XnoObj :(type = tSegment
93             AND SEGinstanceInfo(oSeid).gl.sharable = TRUE)
94 OR (type INSET {tFile, tDevice, tSubtype, tTerminal, tExtent}
95     AND (EXISTS seid pSeidl; openDescriptor od :
96         FCAopenEntry(pSeidl, od).openSeid = oSeid));
97
98     EFFECTS
99     EFFECTS_OF TIIsetEntityLevel(pSeid, oSeid, otii);
100
101 END_MODULE

```

```

1 $(" MODULE:      mac.specs (version 4.5)
2   CONTENTS:      Machine
3   TYPE:          SPECIAL.specifications
4   LAST CHANGED:  7/28/80, 09:52:15
5   ")
6
7
8 MODULE mac
9
10  PARAMETERS
11
12  INTEGER MACmaxVAddr, $( maximum virtual address, also maximum segment size
13                        2^16 * 1 on PDP^11/70)
14  MACmaxOffset, $( maximum offset component of virtual address,
15                  2^13 * 1 on PDP^11/70)
16  MACmaxReg; $( maximum memory mapping register address, seven on
17                PDP^11/70)
18
19  FUNCTIONS
20
21  VFUN MACclock() ^> INTEGER time;
22    $( integer that represents real time)
23    INITIALLY TRUE;
24
25  OFUN MACclockIncrement();
26    $( invlkd continuously by a separate abstract process ^^ the system clock)
27  EFFECTS
28    'MACclock() = MACclock() + 1;
29
30
31 END_MODULE

```

```

1  $("    MODULE:      pbl.specs (version 4.7)
2         CONTENTS:    Parameter Block Functions
3         TYPE:        SPECIAL specifications
4         LAST CHANGED: 12/23/80, 17:17:46
5  ")
6
7
8  MODULE pbl
9
10 $( This module specifies the action of getting arguments or putting results
11    of operations into the virtual memory of a process. The part of
12    virtual memory so manipulated is called a parameter block. The need
13    for parameter blocks comes about when the length of the data is not
14    constant for all invocations of a given operation.
15
16    To make specifications more readable, all parameter block operations are
17    specified in two parts. The basic functionality of an operation is
18    specified in the module to whose object the operation refers, e.g.,
19    the basic specification of readBlock comes from the fmi module. The
20    parameter block manipulation, along with appropriate data conversion,
21    is specified here. This decomposition removes the issue of parameter
22    blocks from the basic specification of already complicated operations.
23    The parameter block manipulation becomes simple once isolated here.)
24
25
26    TYPES
27
28    $(from mac)
29 vAddrType: {0 .. MACmaxVAddr};
30
31    $(from pvm)
32 virtualLocation:
33    STRUCT_OF(domainType domain; spaceType idSpace; vAddrType vAddr);
34 pBlock: STRUCT_OF(virtualLocation vloc; INTEGER size);
35
36    $(from fca)
37 asyncId: CHAR;
38 fileStatus: STRUCT_OF(INTEGER nBlocks, linkCount, timeLastMod; seid subtype;
39    BOOLEAN openAtCrash);
40 ioStatus: STRUCT_OF(INTEGER devDep);
41 fileBlock: VECTOR_OF CHAR;
42 readResult: STRUCT_OF(VECTOR_OF fileBlock data; ioStatus errst);
43
44    $(from spf)
45 SPFargs: VECTOR_OF INTEGER;
46
47
48    EXTERNALREFS
49
50    FROM mac:
51    INTEGER MACmaxVAddr;
52
53    FROM smx:
54    seid: DESIGNATOR;
55    domainType: {nullDomain, kernelDomain, supervisorDomain, userDomain};
56

```

```

57 FROM pvm:
58 spaceType: {nullSpace, iSpace, dSpace};
59 OFUN PVMstore(seid pSeid; pBlock block; VECTOR OF INTEGER vec);
60 VFUN PVMretrieve(seid pSeid; pBlock block) ^> VECTOR_OF INTEGER vec;
61
62 FROM fca:
63 openDescriptor: DESIGNATOR;
64 IOfunction: {rewind, etc};
65
66 FROM fmi:
67 VFUN FCAvDeviceFunction(seid pSeid; openDescriptor od; IOfunction f;
68     VECTOR_OF INTEGER args; asyncId id)
69     ^> ioStatus status;
70 OVFUN FCAdeviceFunction(seid pSeid; openDescriptor od; IOfunction f;
71     VECTOR_OF INTEGER args; asyncId id)
72     ^> ioStatus status;
73 VFUN FCAvReadBlocks(seid pSeid; openDescriptor od; INTEGER blockNo, size;
74     asyncId as) ^> readResult rr;
75 OVFUN FCAreadBlocks(seid pSeid; openDescriptor od; INTEGER blockNo, size;
76     asyncId as) ^> readResult rr;
77 OVFUN FCAwriteBlocks(seid pSeid; openDescriptor od; INTEGER blockNo;
78     VECTOR_OF fileBlock vfb; asyncId id) ^> ioStatus ios;
79
80 FROM spf:
81 SPFFunctionType: {syncSPF, immSegSPF, sysHaltSPF, levelSetSPF};
82 OFUN SPFspecialFunction(seid pSeid; SPFFunctionType fn; SPFargs args);
83
84
85 ASSERTIONS
86
87 FORALL VECTOR OF fileBlock vfb
88 : PBLwordsToBlocks(PBLblocksToWords(vfb)) = vfb;
89
90
91 FUNCTIONS
92
93 $(***** data conversion functions *****)
94
95 VFUN PBLioStatToVec(ioStatus ios) ^> VECTOR_OF INTEGER vi;
96     HIDDEN;
97     INITIALLY vi ~= ?;
98
99 VFUN PBLblocksToWords(VECTOR_OF fileBlock vfb) ^> VECTOR_OF INTEGER vi;
100     HIDDEN;
101     INITIALLY vi ~= ?;
102
103 VFUN PBLwordsToBlocks(VECTOR_OF INTEGER vi) ^> VECTOR_OF fileBlock vfb;
104     HIDDEN;
105     DERIVATION SOME VECTOR_OF fileBlock vfb1 | PBLblocksToWords(vfb1) = vi;
106
107 $(***** operations *****)
108
109 OFUN PBLdeviceFunction(seid pSeid; openDescriptor od; IOfunction f;
110     pBlock arguments, status; asyncId id);
111     $(PBLdeviceFunction)
112
113 DEFINITIONS

```

```

113 VECTOR OF INTEGER inargs IS PVMretrieve(pSeid, arguments);
114 ioStatus st IS FCAvDeviceFunction(pSeid, od, f, inargs, id);
115 VECTOR OF INTEGER result IS PBLioStatToVec(st);
116 EXCEPTIONS
117 EXCEPTIONS OF PVMretrieve(pSeid, arguments);
118 EXCEPTIONS OF FCAvDeviceFunction(pSeid, od, f, inargs, id);
119 EXCEPTIONS OF PVMstore(pSeid, status, result);
120 EFFECTS
121 EFFECTS OF PVMstore(pSeid, status, result);
122 st = EFFECTS OF FCAvDeviceFunction(pSeid, od, f, inargs, id);
123
124 OVFUN PBLreadBlock(seid pSeid; openDescriptor od; INTEGER blockNo;
125 pBlock duFile; asyncId as)
126 *> STRUCT OF (INTEGER bytesRead; ioStatus errst) result; $(PBLreadBlock)
127 DEFINITIONS
128 readResult rr IS FCAvReadBlocks(pSeid, od, blockNo, duFile.size, as);
129 VECTOR OF INTEGER intData IS PBLblocksToWords(rr.data);
130 EXCEPTIONS
131 EXCEPTIONS OF FCAvReadBlocks(pSeid, od, blockNo, duFile.size, as);
132 EXCEPTIONS OF PVMstore(pSeid, duFile, intData);
133 EFFECTS
134 result = STRUCT(LENGTH(intData), rr.errst);
135 EFFECTS OF PVMstore(pSeid, duFile, intData);
136
137 OVFUN PBLspecialFunction(seid pSeid; SPFFunctionType fn; pBlock parm);
138 $(PBLspecialFunction)
139 DEFINITIONS
140 SPFargs args IS PVMretrieve(pSeid, parm);
141 EXCEPTIONS
142 EXCEPTIONS OF PVMretrieve(pSeid, parm);
143 EXCEPTIONS OF SPFspecialFunction(pSeid, fn, args);
144 EFFECTS
145 EFFECTS OF SPFspecialFunction(pSeid, fn, args);
146
147 OVFUN PBLwriteBlock(seid pSeid; openDescriptor od; INTEGER blockNo;
148 pBlock duFile; asyncId id) *> ioStatus ios;
149 $(PBLwriteBlock)
150 DEFINITIONS
151 VECTOR OF fileBlock vfb IS PBLwordsToBlocks(PVMretrieve(pSeid, duFile));
152 EXCEPTIONS
153 EXCEPTIONS OF PVMretrieve(pSeid, duFile);
154 EXCEPTIONS OF FCAwriteBlocks(pSeid, od, blockNo, vfb, id);
155 EFFECTS
156 ios = EFFECTS OF FCAwriteBlocks(pSeid, od, blockNo, vfb, id);
157
158 END_MODULE

```

```

1  $( "      MODULE:      pro.specs (version 4.11)
2          CONTENTS:      Process Operators
3          TYPE:          SPECIAL.specifications
4          LAST CHANGED:   12/23/80, 16:56:41
5  " )
6
7
8  MODULE pro
9
10 $( this module now contains the material which was formerly in
11    the pro, ipc and pst modules)
12
13    TYPES
14
15    $(types supporting pseudo interrupts)
16    piLevelType: {PROminPiLevel .. PROmaxPiLevel}; $(pseudo interrupt level range)
17    piEntryType: STRUCT_OF(piLevelType oldPil;
18                          INTEGER oldPc;
19                          INTEGER oldPs;
20                          ipcMessageType parameter;
21                          INTEGER newPc;
22                          INTEGER newPs);
23    piVectorType: {VECTOR_OF piEntryType piv |
24                  LENGTH(piv) = PROmaxPiLevel-PROminPiLevel+1};
25
26    $(types supporting ipc)
27    ipcqType: {VECTOR_OF ipcMessageType zz | LENGTH(zz) <= IPCmaxMessageCount};
28    ipcTextType: {VECTOR_OF CHAR vc | LENGTH(vc) = IPCmaxMessageLength};
29    ipcMessageType: STRUCT_OF(seid sender; ipcTextType text);
30    pendingType: STRUCT_OF(BOOLEAN flag; INTEGER time);
31
32    $(structure of process status information)
33    processStatusType: STRUCT_OF(seid self; $(part of the process state that users can state)
34                                seid parent;
35                                INTEGER family;
36                                INTEGER realUser;
37                                INTEGER realGroup;
38                                INTEGER advPrio;          $(advisory priority)
39                                piLevelType pil;
40                                INTEGER ps;
41                                INTEGER timerAlarm;
42                                INTEGER clock;
43                                BOOLEAN timTog;
44                                INTEGER supervisorTiming;
45                                INTEGER userTiming);
46    processStateType: STRUCT_OF(INTEGER pc; $(program counter)
47                                VECTOR_OF BOOLEAN vpend;
48                                piVectorType piv;
49                                ipcqType ipcq;
50                                processStatusType pStat);
51
52    $(from smx)
53    nonDisType: STRUCT_OF(
54                      INTEGER securityLevel; SET_OF securityCat securityCatS;
55                      INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
56    daType: SET_OF daMode;
57    modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);

```



```

57 ttiStruct: STRUCT_OF(nonDisType nd;
58     modeStruct ms; INTEGER owner, group; SET_OF privType priv);
59
60
61     PARAMETERS
62
63     INTEGER PROmaxProcessCount;
64
65     INTEGER PROMinPiLevel; $(most interruptable pseudo interrupt level)
66     INTEGER PROmaxPiLevel; $(least interruptable pseudo interrupt level)
67     INTEGER piZero, piIPC, piTimer, piSignal, piIOC, piHardwareFault;
68     $(defined pseudo interrupt levels, I/O completion pseudo*interrupt)
69
70     INTEGER IPCmaxMessageCount; $(maximum number of ipc messages per process)
71     INTEGER IPCmaxMessageLength; $(maximum number of characters per ipc message)
72     INTEGER piPc, $(program counter for process invocation and spawning=0)
73     piPs; $(processor state for process invocation and spawning,
74         154000 octal = user previous domain, supervisor domain, alternate register
75
76     seid processExampleSeid; $(any seid with the tProcess nsp)
77     INTEGER SignalTimeOut;
78
79     DEFINITIONS
80
81     BOOLEAN processExists(seid pSeid) IS PSTprocessState(pSeid) ~= ?;
82     seid newProcessSeid IS $(the process seid generation algorithm)
83     SENmakeSeid(processExampleSeid, SENseidIndex(
84         SOME seid s |
85             SENseidType(s) = tProcess
86             AND NOT processExists(s)));
87     INTEGER processCount IS
88     CARDINALITY({INTEGER i | PSTprocessSlot(i) ~= ?});
89     INTEGER emptySlot IS SOME INTEGER i | i INSET {1 .. PROmaxProcessCount}
90         AND PSTprocessSlot(i) = ?;
91
92
93     EXTERNALREFS
94
95     FROM mac:
96     VFUN MACclock() ^> INTEGER time;
97
98     FROM smx:
99     seid: DESIGNATOR;
100     secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment,
101         tSubtype, tExtent, tNull};
102     privType: {
103         privFileUpdateStatus, privLink,
104         privLockSeg, privModifyPriv,
105         privMount, privSetLevel,
106         privStickySeg, privSetPath,
107         privViolSimpSecurity, privViolStarSecurity,
108         privViolSimpIntegrity, privViolStarIntegrity,
109         privViolDiscrAccess, privRealizeExecPermission,
110         privSignal, privWalkPTable,
111         privHalt, privKernelCall,
112         privViolCompartments, privSetComm,

```

```

113             privImmigrate
114             };
115
116 daMode: {daRead, daWrite, daExecute};
117 securityCat: DESIGNATOR;
118 integrityCat: DESIGNATOR;
119 VFUN SENSEidIndex(seid s) ^> INTEGER index;
120 VFUN SENSEidType(seid s) ^> secureEntityType t;
121 VFUN SENmakeSeid(seid exampleSeid; INTEGER index) ^> seid rSeid;
122 VFUN SMXhasPriv(seid pSeid; privType priv) ^> BOOLEAN b;
123 VFUN SMXflow(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
124 VFUN SMXdap(seid pSeid, oSeid, daType d) ^> BOOLEAN b;
125 VFUN TIigetEntityLevel(seid pSeid, oSeid) ^> tiiStruct otii;
126 VFUN TIInfo(seid s) ^> tiiStruct tiis;
127 OFUN TIisetEntityLevel(seid pSeid, oSied; tiiStruct ntii);
128
129 FROM pvm:
130 segDes: DESIGNATOR;
131
132 FROM pvp:
133 OFUN PVPforkSupport(seid parent, child);
134 OFUN PVPinvokeSupport(seid pSeid, immSeid; segDes arg);
135 OFUN PVPspawnSupport(seid parent, child, imSeid; segDes arg);
136 OFUN PVPreleaseProcessSupport(seid pSeid);
137
138 FROM fca:
139 OFUN FCAcreateOpenTable(seid pSeid);           $(spawn support)
140
141 FROM fmi:
142 OFUN FMIforkSupport(seid parent, child);
143 OFUN FMireleaseSupport(seid pSeid);
144
145
146 ASSERTIONS
147 PROMinPiLevel <= piZero; piZero < piIPC; piIPC < piTimer; piTimer < piSignal; piSignal < p:
148 piIOC < piHardwareFault; piHardwareFault <= PROMaxPiLevel;
149 $( ordering of defined pseudo interrupt levels)
150 SENSEidType(processExampleSeid) = tProcess;
151 $(processExampleSeid is an example of a process seid)
152 processCount <= PROMaxProcessCount; $(there are never too many processes)
153 FORALL seid s: SMXflow(s, processExampleSeid, {daRead});
154 $(the processExampleSeid can be referenced from any security level)
155
156
157 FUNCTIONS
158
159 $( ***** process state functions ***** )
160
161 VFUN PSTprocessState(seid pSeid) ^> processStateType ps;
162     HIDDEN;
163     INITIALLY ps = ?;
164
165 VFUN PSTprocessStatus(seid pSeid) ^> processStatusType ps;
166     HIDDEN;
167     DERIVATION PSTprocessState(pSeid).pStat;
168

```

```

169 VFUN PSTprocessSlot(INTEGER n) ^> seid ps;
170     HIDDEN;
171     INITIALLY ps = ?;
172
173 $( *****support for K_walk_process_table ***** )
174
175 VFUN PROwalkProcessTable(seid pSeid; INTEGER n) ^> seid rSeid;
176     EXCEPTIONS
177         XnoPriv: NOT privWalkPTable INSET TIIinfo(pSeid).priv;
178         XbadSlot: NOT n+1 INSET {1 .. PROmaxProcessCount};
179         Xempty: PSTprocessSlot(n) = ?;
180     DERIVATION
181         PSTprocessSlot(n);
182
183 $( ***** support for K_nap ***** )
184
185 VFUN PSTtimeOfNap(seid pSeid) ^> INTEGER time;
186     HIDDEN;
187     INITIALLY time = 0;
188
189 OFUN PROnap(seid pSeid; INTEGER timeout);
190     DELAY WITH 'PSTtimeOfNap(pSeid) = MACclock();
191             UNTIL MACclock() >= PSTtimeOfNap(pSeid) + timeout
192                   OR PSTprocessStatus(pSeid).pil > piIPC;
193     $(Wakeup on pseudo interrupt occurs when it becomes pending.
194     After waking, the pseudo interrupt is asserted.)
195
196 $( ****support for pseudo interrupts: K_signal and K_interruptReturn**** )
197
198 VFUN PSTtimeOfSignal(seid pSeid) ^> INTEGER time;
199     HIDDEN;
200     INITIALLY time = 0;
201
202 OFUN PROsignal(seid sender, receiver; ipcTextType signalMsg);
203     DEFINITIONS
204         piEntryType receiversSignalEntry()
205             IS PSTprocessState(receiver).piv[piSignal];
206     EXCEPTIONS
207         XnoPriv: NOT SMXhasPriv(sender, privSignal);
208         XbPrSd: NOT (processExists(receiver)
209                     AND SMXflow(sender, receiver, {daRead, daWrite}));
210         XunInt:
211             NOT PSTprocessState(receiver).pStat.pil > piSignal;
212     DELAY WITH 'PSTtimeOfSignal(sender) = MACclock();
213     UNTIL MACclock() >= PSTtimeOfSignal(sender) + SignalTimeOut
214           OR PSTprocessState(receiver).pStat.pil <= piSignal;
215     ASSERTIONS
216         processExists(sender);
217     EFFECTS
218         'PSTprocessState(receiver).vpend[piSignal] = TRUE;
219         'PSTprocessState(receiver).piv[piSignal].parameter.text = signalMsg;
220         'PSTprocessState(receiver).piv[piSignal].parameter.sender = sender;
221 $(PROBLEMS
222     Design requires that K_signal interrupt long kernel calls. How is this
223     this to be done?)
224     $("the only Kcalls which will be interrupted are K^n timer,

```

```

225         K_receive (time out) and a synchronous K_read_block
226         from the terminal. At least for the K_nap and K_receive
227         delays, this can be specified via another condition in the
228         PROnap and PROreceive")
229
230 OFUN PROInterruptReturn(seid pSeid);
231 $(emulates rti instruction. Uses pseudo interrupt vector entry associated
232 with current level and restores pc, ps, and pil to their "old" values)
233 DEFINITIONS
234     piEntryType currentInterruptEntry
235     IS PSTprocessState(pSeid).piv[PSTprocessState(pSeid).pStat.pil];
236 EFFECTS
237     'PSTprocessState(pSeid).pc = currentInterruptEntry.oldPc;
238     'PSTprocessState(pSeid).pStat.ps = currentInterruptEntry.oldPs;
239     'PSTprocessState(pSeid).pStat.pil = currentInterruptEntry.oldPil;
240
241 $( ***** support for ipc: K_post and K_receive ***** )
242
243 OFUN PROpost(seid sender,receiver; BOOLEAN postPseudoInterrupt;
244             ipcTextType text);
245 $(Detects security violation and overflow. Appends message to the tail of
246 the receivers ipcq. Posts pseudo interrupt in the receiver if requested
247 and if receiver has no pending receive)
248 DEFINITIONS
249     VECTOR OF ipcMessageType queue() IS PSTprocessState(receiver).ipcq;
250     INTEGER qLength IS LENGTH(queue());
251 $(There ARE exceptions in this function, but writing up is permitted
252 and the presence or absence of an exception might give information about
253 the state of a write-only object,NOTE THIS AND FIX IT)
254 ASSERTIONS
255     processExists(sender);
256 EFFECTS
257     processExists(receiver)
258     AND SMXflow(sender, receiver, {daWrite})
259     AND qLength < IPCmaxMessageLength
260     AND SMXdap(sender,receiver, {daWrite})
261 => 'PSTprocessState(receiver).ipcq
262     = VECTOR(FOR i FROM 1 TO MIN({1+qLength, IPCmaxMessageLength})
263             : IF i <= qLength
264                 THEN queue()[i]
265                 ELSE STRUCT(sender, text));
266     $(append new message to receivers queue)
267     (postPseudoInterrupt
268     AND NOT PSTreceivePending(receiver).flag
269     => 'PSTprocessState(receiver).vpending[pilIPC] = TRUE);
270     $(post ipc pseudo interrupt if required and if receiver has
271     no pending read)
272
273 VFUN PSTreceivePending(seid pSeid) ^> pendingType r;
274     HIDDEN;
275     INITIALLY r.flag = FALSE AND r.time = 0;
276
277 OVFUN PROreceive(seid pSeid; INTEGER timeout) ^> ipcMessageType msg;
278 $(Returns the ipc message at the head of the queue if one exists or arrives
279 before the expiration of timeout. Otherwise it returns an exception)
280 DEFINITIONS

```

```

281 VECTOR OF ipcMessageType queue() IS PSTprocessState(pSeid).ipcq;
282 INTEGER qLength IS LENGTH(queue());
283 pendingType pending() IS PSTreceivePending(pSeid);
284 EXCEPTIONS
285 XbNoMes: qLength = 0;
286 DELAY WITH 'PSTreceivePending(pSeid).flag = TRUE;
287 'PSTreceivePending(pSeid).time = MACclock();
288 UNTIL qLength > 0 OR MACclock() >= pending().time + timeout;
289 EFFECTS
290 IF qLength > 0 THEN msg = queue()[1] ELSE TRUE;
291 'PSTprocessState(pSeid).ipcq = VECTOR( FOR i FROM 1 TO qLength*1 : queue()[i+1]);
292 'PSTreceivePending(pSeid).flag = FALSE;
293
294 $( ***** support for K_fork ***** )
295
296 OVFUN PROfork(seid parent) *> seid child;
297 DEFINITIONS
298 processStateType p IS PSTprocessState(parent);
299 seid c IS newProcessSeid; processStatusType s IS p.pStat;
300 EXCEPTIONS
301 XnoProc: processCount+1 > PROmaxProcessCount;
302 EXCEPTIONS_OF FMiforkSupport(parent, c);
303 EXCEPTIONS_OF PVPforkSupport(parent, c);
304 EFFECTS
305 child = c;
306 'PSTprocessSlot(emptySlot) = c;
307 'PSTprocessState(c) = STRUCT(p.pc,
308                                VECTOR(),
309                                VECTOR(),
310                                p.ipcq,
311                                STRUCT(c,
312                                        parent,
313                                        s.family,
314                                        s.realUser,
315                                        s.realGroup,
316                                        s.advPrio,
317                                        PROmaxPiLevel,
318                                        s.ps,
319                                        0,
320                                        ?,
321                                        FALSE,
322                                        s.supervisorTiming,
323                                        s.userTiming));
324 $(this assertion specifies the initial tdi state of a forked child)
325 'TIIinfo(c) = TIIinfo(parent);
326 $(this assertion specifies the initial tii state of a forked child)
327 EFFECTS_OF FMiforkSupport(parent, c); $(provide and copy pofv)
328 EFFECTS_OF PVPforkSupport(parent, c); $(provide virtual memory)
329
330 $( ***** support for K_invoke ***** )
331
332 OFUN PROinvoke(seid pSeid, immSeid; segDes arg);
333 DEFINITIONS
334 processStateType p IS PSTprocessState(pSeid);
335 processStatusType s IS p.pStat;
336 tiiStruct pt IS TIIinfo(pSeid);

```

```

337 EXCEPTIONS
338 EXCEPTIONS_OF PVPinvokeSupport(pSeid, immSeid, arg);
339 ASSERTIONS
340 processExists(pSeid);
341 EFFECTS
342 'PSTprocessState(pSeid) = STRUCT(piPc,
343                                VECTOR(),
344                                VECTOR(),
345                                VECTOR(),
346                                STRUCT(s.self,
347                                      s.parent,
348                                      s.family,
349                                      s.realUser,
350                                      s.realGroup,
351                                      s.advPrio,
352                                      PROmaxPiLevel,
353                                      piPs,
354                                      0,
355                                      ?,
356                                      FALSE,
357                                      0,
358                                      0));
359 $(assertion defines the initial process state of the invoked intermediary)
360 'TIIinfo(pSeid) = STRUCT(pt.nd, pt.ms, pt.owner, pt.group,
361                          TIIinfo(immSeid).priv);
362 $(this is how the post^invoke process gets the intermeds privileges)
363 EFFECTS_OF PVPinvokeSupport(pSeid, immSeid, arg); $(redo virtual memory)
364
365 $( "***** support for K_spawn *****")
366
367 OVFUN PROspawn(seid parent, immSeid; segDes arg) ^> seid child;
368 DEFINITIONS
369 processStateType p IS PSTprocessState(parent);
370 processStatusType s IS p.pStat;
371 tiiStruct pt IS TIIinfo(parent);
372 seid c IS newProcessSeid;
373 EXCEPTIONS
374 XnoProc: processCount+1 > PROmaxProcessCount;
375 EXCEPTIONS_OF PVPspawnSupport(parent, c, immSeid, arg);
376 EXCEPTIONS_OF FCACreateOpenTable(parent);
377 ASSERTIONS
378 processExists(parent);
379 EFFECTS
380 child = c;
381 'PSTprocessSlot(emptySlot) = c;
382 'PSTprocessState(c) = STRUCT(piPc,
383                                VECTOR(),
384                                VECTOR(),
385                                VECTOR(),
386                                STRUCT(c,
387                                      s.parent,
388                                      s.family,
389                                      s.realUser,
390                                      s.realGroup,
391                                      s.advPrio,
392                                      PROmaxPiLevel,

```

```

393                                     piPs,
394                                     0,
395                                     ?,
396                                     FALSE,
397                                     0,
398                                     0));
399     $(the process state of the newly spawned intermediary)
400     'TIIinfo(c) = STRUCT(pt.nd, pt.ms, pt.owner, pt.group,
401                          TIIinfo(immSeid).priv);
402     $(post*spawn child acquires intermediatarys privileges)
403     EFFECTS_OF PVPspawnSupport(parent, c, immSeid, arg); $(create vm)
404     EFFECTS_OF FCAcreateOpenTable(c); $(create pofv)
405
406     $( ***** support for K_release_process ***** )
407
408     OFUN PROreleaseProcess(seid pSeid, rSeid);
409     $(Typically a process will release itself and pSeid=rSeid. However this
410     is not treated as a special case.)
411     EXCEPTIONS
412         XbPrSd: (NOT processExists(rSeid))
413                 OR (NOT SMXflow(pSeid, rSeid, {daRead, daWrite}))
414                 OR ((TIIinfo(rSeid).owner ~= TIIinfo(pSeid).owner))
415                     AND NOT SMXhasPriv(pSeid, privSetLevel);
416     ASSERTIONS
417         processExists(pSeid);
418     EFFECTS
419         'PSTprocessSlot(SOME INTEGER i | PSTprocessSlot(i) = rSeid) = ?;
420         'PSTprocessState(rSeid) = ?;
421         'TIIinfo(rSeid) = ?;
422         EFFECTS_OF FMireleaseSupport(rSeid);
423         EFFECTS_OF PVPreleaseProcessSupport(rSeid);
424
425     $( ***** status getting and setting ***** )
426
427     VFUN PROgetProcessStatus(seid pSeid, oSeid) ^> processStatusType ps;
428     DEFINITIONS
429         processStatusType s IS PSTprocessStatus(oSeid);
430     EXCEPTIONS
431         XbPrSd: NOT (processExists(oSeid)
432                     AND SMXflow(pSeid, oSeid, {daRead}))
433                 AND SMXdap(pSeid, oSeid, {daRead}));
434     ASSERTIONS
435         processExists(pSeid);
436     DERIVATION
437         STRUCT(s.self,
438               s.parent,
439               s.family,
440               s.realUser,
441               s.realGroup,
442               s.advPrio,
443               s.pil,
444               s.ps,
445               s.timerAlarm,
446               MACclock(),
447               s.timTog,
448               s.supervisorTiming,

```

```

449         s.userTiming);
450
451 OFUN PROsetProcessStatus(seid pSeid, oSeid; processStatusType n);
452 DEFINITIONS
453     processStateType e IS PSTprocessState(oSeid);
454     processStatusType o IS PSTprocessStatus(oSeid);
455 EXCEPTIONS
456     XbPrSd: (NOT processExists(oSeid))
457         OR (NOT SMXflow(pSeid, oSeid, {daRead, daWrite}))
458         OR ((TIIinfo(oSeid).owner ~= TIIinfo(pSeid).owner))
459         AND NOT SMXhasPriv(pSeid, privSetLevel);
460 EXCEPTIONS OF PROreleaseProcess(pSeid,oSeid);
461 ASSERTIONS
462     processExists(pSeid);
463 EFFECTS
464     'PSTprocessState(oSeid) = STRUCT(e.pc,
465                                     e.vpend,
466                                     e.piv,
467                                     e.ipcq,
468                                     STRUCT(o.self,
469                                         o.parent,
470                                         n.family,
471                                         n.realUser,
472                                         n.realGroup,
473                                         n.advPrio,
474                                         n.pil,
475                                         n.ps,
476                                         n.timerAlarm,
477                                         ?,
478                                         n.timTog,
479                                         o.supervisorTiming,
480                                         o.userTiming));
481 END_MODULE

```



```

1  $(  
2      MODULE:      pvm.specs (version 4.11)  
3      CONTENTS:    Virtual Memory  
4      TYPE:        SPECIAL.specifications  
5      LAST CHANGED: 12/23/80, 17:37:23  
6  ")  
7  
8  MODULE pvm  
9  
10 $( this module now contains the contents of what was the seg and pvm modules)  
11  
12     TYPES  
13  
14     $(FROM mac)  
15 vAddrType: {0 .. MACmaxVaddr};  
16     $(FROM smx "" tii)  
17 daType: SET_OF daMode;  
18 modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);  
19 domainType: {nullDomain, kernelDomain, supervisorDomain, userDomain};  
20 nonDisType: STRUCT_OF(  
21     INTEGER securityLevel; SET_OF securityCat securityCatS;  
22     INTEGER integrityLevel; SET_OF integrityCat integrityCatS);  
23 tiiStruct: STRUCT_OF(nonDisType nd;  
24     modeStruct ms; INTEGER owner, group; SET_OF privType priv);  
25     $(from seg "" exportable)  
26 segDes: DESIGNATOR;  
27 spaceType: {nullSpace, iSpace, dSpace};  
28 direction: {up, down};  
29     $(from seg "" redeclarable)  
30 virtualLocation:  
31     STRUCT_OF(domainType domain; spaceType idSpace; vAddrType vAddr);  
32     $(from seg "" local)  
33 globalData:  
34     STRUCT_OF(BOOLEAN sharable, swappable, sticky, memAdvise, executable);  
35 instanceStruct:  
36     STRUCT_OF(globalData gl; direction growth; INTEGER refCount;  
37     VECTOR_OF INTEGER data);  
38 useStruct: STRUCT_OF(seid instance; virtualLocation vloc; daType da);  
39     $(from pvm "" redeclarable)  
40 pBlock: STRUCT_OF(virtualLocation vloc; INTEGER size);  
41 statusStruct:  
42     STRUCT_OF(globalData gl; direction growth; INTEGER size; seid iseid;  
43     segDes udes; virtualLocation vloc; daType da; BOOLEAN mapped);  
44  
45  
46     PARAMETERS  
47  
48     $(from seg)seid exampleSegmentSeid; $(used for segment creation)  
49 segDes SEGnullSeg; $( indicates null segment designator)  
50 INTEGER PVMmaxSegDes; $(maximum number of segment designators in an address  
51     space)  
52  
53  
54     DEFINITIONS  
55  
56     $(from seg)

```

```

57 INTEGER SEGsize(seid segSeid) IS LENGTH(SEGinstanceInfo(segSeid).data);
58
59 $(from pvm)
60 INTEGER nSegs(seid pSeid)
61 IS CARDINALITY({segDes sd | SEGuseInfo(pSeid, sd) ~= ?});
62
63 BOOLEAN PVMvmExists(seid pSeid) IS PVMsegmentSet(pSeid) ~= ?;
64
65 segDes PVMblockToSeg(seid pSeid; pBlock block) IS
66 SOME segDes sd
67 | sd INSET PVMmappedSegmentSet(pSeid)
68 AND (EXISTS useStruct use = SEGuseInfo(pSeid, sd)
69 : use.vloc.domain = block.vloc.domain
70 AND use.vloc.idSpace = block.vloc.idSpace
71 AND use.vloc.vAddr <= block.vloc.vAddr
72 AND block.vloc.vAddr + block.size
73 <= use.vloc.vAddr + SEGsize(use.instance));
74 $( gives the segment designator, if any, that totally contains block in
75 address space designated by pSeid; if there is none, returns ?; the
76 segment must be mapped)
77
78 SET OF INTEGER addrRegRange(INTEGER vAddr, size; direction d) IS
79 IF d = up
80 THEN {vAddr / MACmaxOffset .. (vAddr + size # 1) / MACmaxOffset}
81 ELSE {(vAddr # size + 1) / MACmaxOffset .. vAddr / MACmaxOffset};
82 $( gives the range of address registers used by a segment as a function of
83 its start address, size, and growth direction)
84
85 SET OF INTEGER addrRegRangeSeg(seid pSeid; segDes s) IS
86 LET useStruct use = SEGuseInfo(pSeid, s)
87 IN addrRegRange(use.vloc.vAddr, SEGsize(use.instance),
88 SEGinstanceInfo(use.instance).growth);
89 $( gives the range of address registers used by a segment as a function of
90 the process id and the segment designator)
91
92 BOOLEAN noHole(seid pSeid; INTEGER size; virtualLocation vl; direction d;
93 SET OF segDes ssd) IS
94 NOT addrRegRange(vl.vAddr, size, d) SUBSET {0 .. MACmaxReg}
95 OR (EXISTS segDes s | s INSET ssd; useStruct use = SEGuseInfo(pSeid, s)
96 : use.vloc.idSpace = vl.idSpace
97 AND use.vloc.domain = vl.domain
98 AND addrRegRangeSeg(pSeid, s)
99 INTER addrRegRange(vl.vAddr, size, d) ~= {});
100 $(TRUE iff a segment described by size, vl, and direction will NOT fit into
101 a hole in the address space designated by pSeid and ssd; this includes
102 testing for virtual memory underflow and overflow)
103
104 EXTERNALREFS
105
106 FROM mac:
107 INTEGER MACmaxVaddr, MACmaxOffset, MACmaxReg;
108
109 FROM smx:
110 $(FROM sen:)
111 seid: DESIGNATOR;
112 secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,

```

```

113         tExtent, tNull};
114 VFUN SENseidNsp(seid s) '> INTEGER nsp;
115 VFUN SENseidIndex(seid s) '> INTEGER index;
116 VFUN SENmakeSeid(seid exampleSeid; INTEGER index) '> seid outSeid;
117 VFUN SENseidType(seid s) '> secureEntityType set;
118     $(FROM prv:)
119     privType: {
120         privFileUpdateStatus,  privLink,
121         privLockSeg,           privModifyPriv,
122         privMount,             privSetLevel,
123         privStickySeg,         privSetPath,
124         privViolSimpSecurity,  privViolStarSecurity,
125         privViolSimpIntegrity, privViolStarIntegrity,
126         privViolDiscrAccess,  privRealizeExecPermission,
127         privSignal,           privWalkPTable,
128         privHalt,             privKernelCall,
129         privViolCompartments, privSetComm,
130         privImmigrate
131     };
132     $(FROM tii:)
133     securityCat: DESIGNATOR;
134     integrityCat: DESIGNATOR;
135     daMode: {daRead, daWrite, daExecute};
136 VFUN TIIinfo(seid anySeid) '> tiiStruct tiiSt;
137     $(FROM smx:)
138 VFUN SMXhasPriv(seid pSeid; privType priv) '> BOOLEAN b;
139 VFUN SMXflow(seid pSeid, oSeid; daType d) '> BOOLEAN b;
140 VFUN SMXdap(seid pSeid, oSeid; daType d) '> BOOLEAN b;
141
142
143     ASSERTIONS
144     $(from seg)
145     MACmaxVaddr > 0; MACmaxOffset > 0; MACmaxReg > 0;
146     MACmaxVaddr + 1 = (MACmaxOffset + 1) * (MACmaxReg + 1);
147     PVMmaxSegDes >= 2;
148     $( basic relations among the SEG parameters)
149     SENseidType(exampleSegmentSeid) = tSegment;
150     $( basic property of exampleSegmentSeid and all segment seids)
151     FORALL seid s | SEGinstanceInfo(s) ~= ?
152         : SENseidNsp(s) = SENseidNsp(exampleSegmentSeid);
153     $(all seids for existing segments have a distinbuished nsp component)
154     FORALL seid s | SEGinstanceInfo(s) ~= ? : TIIinfo(s) ~= ?;
155     $(all existing segments have an exiting TII entry)
156     FORALL seid pSeid; segDes segd | SEGuseInfo(pSeid, segd) ~= ?
157         : SEGinstanceInfo(SEGuseInfo(pSeid, segd).instance) ~= ?;
158     $(all valid segment uses have corresponding valid segment instances)
159
160     $(from pvm)
161     FORALL seid pSeid | PVMvmExists(pSeid); segDes sd
162         : (sd INSET PVMsegmentSet(pSeid)) = (SEGuseInfo(pSeid, sd) ~= ?);
163     $(defines what it means for a segment to be in the segment set of a
164     process)
165     FORALL seid pSeid | PVMvmExists(pSeid)
166         : PVMmappedSegmentSet(pSeid) SUBSET PVMsegmentSet(pSeid);
167     $(only existing segments can be mapped)
168     FORALL seid pSeid1, pSeid2; segDes sd1, sd2:

```

```

169 NOT SEGinstanceInfo(SEGuseInfo(pSeid1,sd1).instance).gl.sharable
170 =>
171 (SEGuseInfo(pSeid1,sd1).instance ~= SEGuseInfo(pSeid2,sd2).instance);
172 $(A nonsharable instance struct cannot be pointed to by more than 1
173 use struct. Checked for in PVMrendezvous.)
174 FORALL seid pSeid | PVMvmExists(pSeid); segDes s1; segDes s2
175 : LET useStruct use1 = SEGuseInfo(pSeid, s1);
176 useStruct use2 = SEGuseInfo(pSeid, s2)
177 IN s1 ~= s2 AND use1 ~= ? AND use2 ~= ?
178 AND {s1, s2} SUBSET PVMmappedSegmentSet(pSeid)
179 AND use1.vloc.domain = use2.vloc.domain
180 AND use1.vloc.idSpace = use2.vloc.idSpace
181 => addrRegRangeSeg(pSeid, s1) INTER addrRegRangeSeg(pSeid, s2) = {};
182 $(no two mapped segments in the same domain and idSpace may have
183 overlapping memory address registers)
184
185
186 FUNCTIONS
187
188 $(***** state functions *****)
189
190 VFUN SEGinUseIndexSet() ^> SET_OF INTEGER s1; $(SEGinUseIndexSet)
191 $( gives the set of seid indices that are currently in use for existing
192 segments)
193 HIDDEN;
194 INITIALLY s1 = {};
195
196 VFUN SEGinstanceInfo(seid segSeid) ^> instanceStruct is; $(SEGinstanceInfo)
197 $( gives all the information pertaining to a segment's global data,
198 referred to as segment in tance data)
199 HIDDEN;
200 INITIALLY is = ?;
201
202 VFUN SEGuseInfo(seid pSeid; segDes segd) ^> useStruct us; $(SEGuseInfo)
203 $( gives all the information pertaining to a segment's use in the address
204 space in a particular process; this is information local to a process)
205 HIDDEN;
206 INITIALLY us = ?;
207
208 VFUN PVMsegmentSet(seid pSeid) ^> SET_OF segDes segSet; $(PVMsegmentSet)
209 $( gives the set of segments possessed by a given process)
210 INITIALLY segSet = ?;
211
212 VFUN PVMmappedSegmentSet(seid pSeid) ^> SET_OF segDes mappedSet;
213 $(PVMmappedSegmentSet)
214 $( gives the set of mapped " or active segments " of a process; a
215 segment cannot be addressed unless it is mapped)
216 HIDDEN;
217 INITIALLY mappedSet = ?;
218
219 $(***** virtual memory management *****)
220
221 CFUN PVMcreateVM(seid pSeid); $(PVMcreateVM)
222 $( creates a new virtual memory to be identified by "pSeid"; this VM
223 must not currently exist)
224 ASSERTIONS

```

```

225 NOT PVMvmExists(pSeid);
226 EFFECTS
227 'PVMsegmentSet(pSeid) = {};
228 'PVMmappedSegmentSet(pSeid) = {};
229
230 OFUN PVMdeleteVM(seid pSeid); $(PVMdeleteVM)
231 $( deletes the currently existing virtual memory "pSeid")
232 ASSERTIONS
233 PVMvmExists(pSeid);
234 EFFECTS
235 'PVMsegmentSet(pSeid) = ?;
236 'PVMmappedSegmentSet(pSeid) = ?;
237
238 $(***** basic segment management *****)
239
240 OFUN PVMstore(seid pSeid; pBlock block; VECTOR_OF INTEGER vec); $(PVMstore)
241 $(inserts contents of vec into the mapped segment indicated by block)
242 DEFINITIONS
243 segDes targ IS PVMblockToSeg(pSeid, block);
244 useStruct use IS SEGuseInfo(pSeid, targ);
245 instanceStruct inst IS SEGinstanceInfo(use.instance);
246 EXCEPTIONS
247 XbAddr: targ = ?;
248 XbadDa: NOT daWrite INSET use.da;
249 XbSgRng: LENGTH(vec) ~= block.size;
250 XbSgSz: block.size <= 0; $(IMPOSSIBLE, CARDINAL)
251 ASSERTIONS
252 PVMvmExists(pSeid);
253 EFFECTS
254 LET INTEGER relOffset = block.vloc.vAddr # use.vloc.vAddr
255 IN 'SEGinstanceInfo(use.instance) =
256 STRUCT(inst.gl, inst.growth, inst.refCount,
257 VECTOR(
258 FOR i FROM 1 TO LENGTH(inst.data)
259 : IF i INSET {relOffset + 1 .. relOffset + LENGTH(vec)}
260 THEN vec[i # relOffset]
261 ELSE inst.data[i]));
262
263 VFUN PVMretrieve(seid pSeid; pBlock block) #> VECTOR_OF INTEGER vec;
264 $( retrieves data from a mapped segment as specified by pBlock)
265 DEFINITIONS
266 segDes targ IS PVMblockToSeg(pSeid, block);
267 useStruct use IS SEGuseInfo(pSeid, targ);
268 instanceStruct inst IS SEGinstanceInfo(use.instance);
269 EXCEPTIONS
270 XbSgSz: block.size <= 0; $(IMPOSSIBLE CARDINAL)
271 XbSgDes: targ = ?;
272 ASSERTIONS
273 PVMvmExists(pSeid);
274 DERIVATION
275 VECTOR(FOR i FROM 1 TO block.size:
276 inst.data[block.vloc.vAddr # use.vloc.vAddr + i # 1]);
277
278 OVFUN PVMbuild(seid pSeid; statusStruct st; modeStruct ms)
279 #> STRUCT_OF(seid segSeid; segDes segd) result; $(PVMbuild)
280 $( builds a new segment with the specified parameters;

```

```

281     an entry in the tii table i also created for the segment;
282     the results are the "" previously unused "" seid for the new segment and
283     the "" previously unused "" segment designator; the newly created
284     segment is mapped, indicating that it can be addressed immediately;
285     discretionary access that the process allows itself to the segment, "m"
286     must be a subset of the owner access specified in the tii information;
287     "ms" of the new segment must be within the size limitations and fit into
288     the mapped virtual memory space of the creating process)
289 DEFINITIONS
290     tiiStruct proTii IS TIIinfo(pSeid);
291     tiiStruct segTii
292     IS STRUCT(proTii.nd, ms, proTii.owner, proTii.group, proTii.priv);
293     seid newSegSeid
294     IS SENmakeSeid(exampleSegmentSeid,
295                     SOME INTEGER i | NOT i INSET SEGinUseIndexSet());
296     segDes sd IS SOME segDes sdl | SEGuseInfo(pSeid, sdl) = ?;
297 EXCEPTIONS
298     XnPvStkSg: st.gl.sticky
299               AND NOT SMXhasPriv(pSeid, privStickySeg);
300     XnPvLkSg: NOT st.gl.swappable AND NOT SMXhasPriv(pSeid, privLockSeg);
301     XbadMode: ((daWrite INSET ms.ownerMode AND NOT
302                daRead INSET ms.ownerMode) OR (daWrite INSET ms.groupMode
303                AND NOT daRead INSET ms.groupMode) OR (daWrite INSET
304                ms.allMode AND NOT daRead INSET ms.allMode));
305     XbSgSz: NOT st.size INSET {1 .. MACmaxVaddr};
306     $(test made only for 0;11/70,s card) $
307     XvrtMcCfl:
308       noHole(pSeid, st.size, st.vloc, st.growth, PVMmappedSegmentSet(pSeid));
309     XpostEh: nSegs(pSeid) >= PVMmaxSegDes;
310     RESOURCE_ERROR; $( ran out of table space or seid space)
311 ASSERTIONS
312     PVMvmExists(pSeid);
313 EFFECTS
314     'TIIinfo(newSegSeid) = segTii;
315     'SEGinUseIndexSet()
316     = SEGinUseIndexSet() UNION {SENseidIndex(newSegSeid)};
317     'SEGinstanceInfo(newSegSeid) =
318     STRUCT(st.gl, st.growth, 1,
319            VECTOR(FOR i FROM 1 TO st.size : 0));
320     'SEGuseInfo(pSeid, sd) = STRUCT(newSegSeid, st.vloc, ms.ownerMode);
321     result = STRUCT(newSegSeid, sd);
322     'PVMsegmentSet(pSeid) = PVMsegmentSet(pSeid) UNION {sd};
323     'PVMmappedSegmentSet(pSeid)
324     = PVMmappedSegmentSet(pSeid) UNION {sd};
325
326 OFUN PVMdestroy(seid pSeid; segDes segd);                                $(PVMdestroy)
327     $(destroys the segment use indicated by pSeid and segd; if the segment is
328     unsticky and otherwise unreferenced, the segment instance information is
329     also deleted)
330 DEFINITIONS
331     useStruct use IS SEGuseInfo(pSeid, segd);
332     seid segSeid IS use.instance;
333     instanceStruct inst IS SEGinstanceInfo(segSeid);
334 EXCEPTIONS
335     KEsegNotHeld: NOT segd INSET PVMsegmentSet(pSeid);
336 ASSERTIONS

```

```

337     PVMvmExists(pSeid);
338 EFFECTS
339     'PVMsegmentSet(pSeid) = PVMsegmentSet(pSeid) DIFF {segd};
340     'PVMmappedSegmentSet(pSeid) = PVMmappedSegmentSet(pSeid) DIFF {segd};
341     'SEGuseInfo(pSeid, segd) = ?;
342     IF (inst.refCount = 1 AND inst.gl.sticky = FALSE)
343     THEN 'SEGinstanceInfo(segSeid) = ?
344         AND 'SEGINuseIndexSet()
345             = SEGINuseIndexSet() DIFF {SENseidIndex(segSeid)}
346         AND 'TIIinfo(segSeid) = ?
347     ELSE 'SEGinstanceInfo(segSeid) =
348         STRUCT(inst.gl, inst.growth, inst.refCount * 1, inst.data);
349
350 $(***** high level storage allocation *****)
351
352 OFUN PVMremap(seid pSeid; segDes in; virtualLocation vl; daType d;
353     BOOLEAN vlFlg, daFlg; segDes out; INTEGER newSize;
354     BOOLEAN nsFlg); $(PVMremap)
355 $( this function takes the currently mapped segment "out" and maps it out,
356     while simultaneously mapping in the currently unmapped segment "in";
357     this function can be used for mapping in "" without mapping out "" by
358     letting "out" be the distinguished value SEGnullSeg; similarly, mapping
359     out alone can be done by letting "in" be SEGnullSeg; a mapped in segment
360     can have a new virtual location and a discretionary access specified
361     optionally; a mapped out segment can have its size optionally changed;
362     all these optional changes are specified by the values of BOOLEAN flags;
363     the idSpace of the mapped in segment may not be changed; the mapped in
364     segment must occupy a hole in the virtual memory)
365 DEFINITIONS
366     SET OF segDes inSet IS IF in = SEGnullSeg THEN {} ELSE {in};
367     SET OF segDes outSet IS IF out = SEGnullSeg THEN {} ELSE {out};
368     useStruct inUse IS SEGuseInfo(pSeid, in);
369     instanceStruct inInst IS SEGinstanceInfo(inUse.instance);
370     useStruct outUse IS SEGuseInfo(pSeid, out);
371     seid outSeid IS outUse.instance;
372     instanceStruct outInst IS SEGinstanceInfo(outSeid);
373 EXCEPTIONS
374     XbSgDes: NOT inSet SUBSET PVMsegmentSet(pSeid);
375     XoutSgAldUmp: NOT outSet SUBSET PVMmappedSegmentSet(pSeid);
376     XwOnlSg: daFlg AND daWrite INSET d AND NOT daRead INSET d;
377     XinSgAldMap: {in} SUBSET PVMmappedSegmentSet(pSeid);
378     XvrtMcCfl:
379         in ~= SEGnullSeg
380         AND noHole(pSeid, SEGsize(inUse.instance),
381             IF vlFlg THEN vl ELSE inUse.vloc, inInst.growth,
382             PVMmappedSegmentSet(pSeid) DIFF outSet);
383     XbNewSz:
384         nsFlg AND (NOT newSize INSET {1 .. MACmaxVaddr}
385             OR outInst.gl.sharable
386             OR NOT daWrite INSET outUse.da);
387     RESOURCE_ERROR; $(for swap space upon segment growth)
388 ASSERTIONS
389     PVMvmExists(pSeid);
390 EFFECTS
391     'PVMmappedSegmentSet(pSeid) = (PVMmappedSegmentSet(pSeid) DIFF outSet)
392         UNION inSet;

```

```

393 'SEGuseInfo(pSeid, in)
394 = STRUCT(inUse.instance, IF vlFlg THEN vl ELSE inUse.vloc,
395          IF daFlg THEN d ELSE inUse.da);
396 (nsFlg AND newSize ~= SEGsize(inUse.instance))
397 => 'SEGinstanceInfo(outSeid)
398 = STRUCT(outInst.gl, outInst.growth, outInst.refCount,
399          VECTOR(FOR i FROM 1 TO newSize
400                : IF i <= SEGsize(outSeid)
401                  THEN outInst.data[i] ELSE 0));
402
403 $(***** segment sharing *****)
404
405 OVFUN PVMrendezvous(seid pSeid, segSeid; virtualLocation vl; daType d)
406   *> segDes sd; $(PVMrendezvous)
407 $( creates a use for the segment named by segSeid; this segment appears
408   inaccessible if the multilevel security model would consider it a
409   violation of information flow)
410 DEFINITIONS
411   instanceStruct inst IS SEGinstanceInfo(segSeid);
412   segDes segd IS SOME segDes sdl | SEGuseInfo(pSeid, sdl) = ?;
413 EXCEPTIONS
414   KSegBadName:
415     NOT (inst ~= ?
416          AND (IF daWrite INSET d
417                THEN SMXflow(pSeid, segSeid, {daRead, daWrite})
418                ELSE SMXflow(pSeid, segSeid, {daRead})));
419   XbadDa: NOT SMXdap(pSeid, segSeid, d);
420   XdupSg: EXISTS segDes sdl
421           : SEGuseInfo(pSeid, sdl).instance = segSeid;
422   XwOnlSg: daWrite INSET d AND NOT daRead INSET d;
423   XvrtMmCfl: noHole(pSeid, SEGsize(segSeid), vl, inst.growth,
424                    PVMmappedSegmentSet(pSeid));
425   XpostEh: nSegs(pSeid) >= PVMmaxSegDes;
426   XnotSh: NOT inst.gl.sharable;
427 ASSERTIONS
428   PVMvmExists(pSeid);
429 EFFECTS
430   'SEGuseInfo(pSeid, segd) = STRUCT(segSeid, vl, d);
431   'SEGinstanceInfo(segSeid)
432   = STRUCT(inst.gl, inst.growth, inst.refCount + 1, inst.data);
433   'PVMsegmentSet(pSeid) = PVMsegmentSet(pSeid) UNION {segd};
434   'PVMmappedSegmentSet(pSeid)
435   = PVMmappedSegmentSet(pSeid) UNION {segd};
436   sd = segd;
437
438 OFUN PVMcopySeg(seid fromSeid, toSeid; segDes sd); $(PVMcopySeg)
439 $( copies a segment from the virtual memory "fromSeid" to the virtual
440   memory "toSeid"; both virtual memories must exist; the segment
441   designator sd must exist in "fromSeid" but not in "toSeid"; used by
442   the module that sets up virtual memories for new processes)
443 DEFINITIONS
444   useStruct use IS SEGuseInfo(fromSeid, sd);
445   seid oldSeid IS use.instance; instanceStruct inst IS SEGinstanceInfo(oldSeid);
446   seid newSeid
447   IS SENmakeSeid(exampleSegmentSeid,
448                  SOME INTEGER i | NOT i INSET SEGinUseIndexSet());

```



```

449     tiiStruct stii IS TIIinfo(oldSeid);
450     tiiStruct ptii IS TIIinfo(toSeid);
451     ASSERTIONS
452         PVMvmExists(fromSeid);
453         PVMvmExists(toSeid);
454         sd INSET PVMsegmentSet(fromSeid);
455         NOT sd INSET PVMsegmentSet(toSeid);
456     EFFECTS
457         'SEGINstanceInfo(newSeid) = STRUCT(inst.gl, inst.growth, 1, inst.data);
458         'SEGuseInfo(toSeid, sd) = STRUCT(newSeid, use.vloc, use.da);
459         'TIIinfo(newSeid) = STRUCT(stii.nd, stii.ms, ptii.owner, ptii.group,
460                                   stii.priv);
461         'PVMsegmentSet(toSeid) = PVMsegmentSet(toSeid) UNION {sd};
462         sd INSET PVMmappedSegmentSet(fromSeid) =>
463         'PVMmappedSegmentSet(toSeid) = PVMmappedSegmentSet(toSeid) UNION {sd};
464         'SEGINuseIndexSet()
465         = SEGINuseIndexSet() UNION {SENseidIndex(newSeid)};
466
467     $(***** segment status manipulation *****)
468
469     VFUN PVMgetSegmentStatus(seid pSeid, segSeid; segDes segd) ^> statusStruct ss;
470                                     $(PVMgetSegmentStatus)
471     $( returns the status information ^^ which is much of the global
472       information ^^ for the segment; the segment must exist in the segment
473       set of the requesting process)
474     DEFINITIONS
475         seid checkSeid IS IF segd=SEGnullSeg THEN segSeid ELSE
476         -                     SEGuseInfo(pSeid,segd).instance;
477         segDes desrefl IS SOME segDes s |
478                     SEGuseInfo(pSeid,s).instance=checkSeid;
479         segDes desref IS IF desrefl = ? THEN SEGnullSeg
480                     ELSE desrefl;
481     EXCEPTIONS
482         XbSgSd: IF segd = SEGnullSeg THEN (SEGINstanceInfo(segSeid) = ?
483                     OR NOT SMXflow(pSeid, segSeid, {daRead}))
484         ELSE NOT segd INSET PVMsegmentSet(pSeid);
485     ASSERTIONS
486         PVMvmExists(pSeid);
487     DERIVATION
488         STRUCT(
489             SEGINstanceInfo(checkSeid).gl,
490             SEGINstanceInfo(checkSeid).growth,
491             SEGsize(checkSeid),
492             checkSeid,
493             desref,
494             IF desref=SEGnullSeg THEN STRUCT(nullDomain,nullSpace,0)
495             ELSE SEGuseInfo(pSeid,desref).vloc,
496             IF desref=SEGnullSeg THEN {}
497             ELSE SEGuseInfo(pSeid,desref).da,
498             IF desref=SEGnullSeg THEN FALSE
499             ELSE IF desref INSET PVMmappedSegmentSet(pSeid) THEN TRUE
500             ELSE FALSE
501         );
502     OFUN PVMsetSegmentStatus(seid pSeid, segSeid; globalData glo);
503                                     $(PVMsetSegmentStatus)
504     $( changes the status information for the segment; certain privileges

```

```

505     are required;)
506 DEFINITIONS
507     instanceStruct i IS SEGinstanceInfo(segSeid);
508 EXCEPTIONS
509     XbSgSd: NOT (i ~= ? AND SMXflow(pSeid, segSeid, {daRead, daWrite}));
510     Xsharable: i.gl.sharable;
511     XswapEh:
512         glo.swappable AND NOT i.gl.swappable
513         AND NOT SMXhasPriv(pSeid, privLockSeg);
514     XnPvStkSg:
515         glo.sticky AND NOT i.gl.sticky
516         AND NOT SMXhasPriv(pSeid, privStickySeg);
517 ASSERTIONS
518     PVMvmExists(pSeid);
519 EFFECTS
520     'SEGinstanceInfo(segSeid) = STRUCT(glo, i.growth, i.refCount, i.data);
521
522 $(***** level getting and setting *****)
523
524 VFUN PVMgetSegmentLevel(seid pSeid, segSeid) '> tiiStruct tii;
525                                     $(PVMgetSegmentLevel)
526 EXCEPTIONS
527     $(LR '' these must be finalized)
528     XbSgSd: NOT (SEGinstanceInfo(segSeid) ~= ?
529                 AND SMXflow(pSeid, segSeid, {daRead}));
530     XbadDa: NOT SMXdap(pSeid, segSeid, {daRead});
531 DERIVATION
532     TIIinfo(segSeid);
533
534 OFUN PVMsetSegmentLevel(seid pSeid, segSeid; tiiStruct ntii);
535                                     $(PVMsetSegmentLevel)
536 DEFINITIONS
537     tiiStruct otii IS TIIinfo(segSeid);
538 EXCEPTIONS
539     XbSgSd: NOT (SEGinstanceInfo(segSeid) ~= ?
540                 AND SMXflow(pSeid, segSeid, {daRead, daWrite}));
541     XbadDa: NOT SMXdap(pSeid, segSeid, {daWrite});
542     XshSg: SEGinstanceInfo(segSeid).gl.sharable;
543     XnoSLev: otii.nd ~= ntii.nd
544                 AND NOT SMXhasPriv(pSeid, privLockSeg);
545     XwOnlSg: otii.ms ~= ntii.ms
546                 AND ((daWrite INSET ntii.ms.ownerMode
547                     AND NOT daRead INSET ntii.ms.ownerMode)
548                     OR (daWrite INSET ntii.ms.groupMode
549                     AND NOT daRead INSET ntii.ms.groupMode)
550                     OR (daWrite INSET ntii.ms.allMode
551                     AND NOT daRead INSET ntii.ms.allMode));
552     XchgSgOwn:
553         (otii.ms ~= ntii.ms OR otii.owner ~= ntii.owner
554          OR otii.group ~= ntii.group)
555         AND otii.owner ~= TIIinfo(pSeid).owner;
556     XnoSPriv: otii.priv ~= ntii.priv
557                 AND NOT SMXhasPriv(pSeid, privModifyPriv);
558 EFFECTS
559     'TIIinfo(segSeid) = ntii;
560

```

561

562 END_MODULE

```

1  $( "    MODULE:      pvp.specs (version 4.12)
2          CONTENTS:    Virtual Memory " Process Supp
3          TYPE:
4          LAST CHANGED: 12/23/80, 16:59:00
5  " )
6
7
8  MODULE pvp
9
10 $( this module contains operations that support the process module's use of
11    the virtual memory mechanism. This is entirely a procedure abstraction.
12    It is specified in terms of V*functions of other modules, but is
13    implemented by a program in terms of operations of other modules.
14    This sequence will be included in the comments for each of the operations.
15    This module has no state of its own except for parameters for immediate
16    segments.)
17
18    TYPES
19
20    $(from mac)
21    vAddrType: {0 .. MACmaxVAddr};
22
23    $(from smx)
24    nonDisType: STRUCT_OF(
25        INTEGER securityLevel; SET_OF securityCat securityCatS;
26        INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
27    daType: SET_OF daMode;
28    modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
29    tiiStruct: STRUCT_OF(nonDisType nd;
30        modeStruct ms; INTEGER owner, group; SET_OF privType priv);
31
32    $(from pvm)
33    virtualLocation:
34        STRUCT_OF(domainType domain; spaceType idSpace; vAddrType vAddr);
35    globalData:
36        STRUCT_OF(BOOLEAN sharable, swappable, sticky, memAdvise, executable;
37            direction growth);
38    instanceStruct:
39        STRUCT_OF(globalData gl; direction growth; INTEGER refCount; VECTOR_OF INTEGER data);
40    useStruct: STRUCT_OF(seid instance; virtualLocation vloc; daType da);
41
42
43    PARAMETERS
44
45    virtualLocation PVMimmVloc; $(location for immediate segment)
46    segDes PVMimmDes; $(designator for immediate text segment)
47    segDes PVMargDes; $(designator for argument segment)
48    virtualLocation PVMargVloc; $(location for argument segment)
49
50
51    DEFINITIONS
52
53    $(from seg)
54    INTEGER SEGsize(seid segSeid) IS LENGTH(SEGinstanceInfo(segSeid).data);
55
56    $(from pvm)

```

```

57 INTEGER nSegs(seid pSeid)
58   IS CARDINALITY({segDes sd | SEGuseInfo(pSeid, sd) ~= ?});
59
60 BOOLEAN PVMvmExists(seid pSeid) IS PVMsegmentSet(pSeid) ~= ?;
61
62 SET OF INTEGER addrRegRange(INTEGER vAddr, size; direction d) IS
63   IF d = 'v'
64     THEN {vAddr / MACmaxOffset .. (vAddr + size * 1) / MACmaxOffset}
65     ELSE {(vAddr * size + 1) / MACmaxOffset .. vAddr / MACmaxOffset};
66   $( gives the range of address registers used by a segment as a function of
67     its start address, size, and growth direction)
68
69 SET OF INTEGER addrRegRangeSeg(seid pSeid; segDes s) IS
70   LET useStruct use = SEGuseInfo(pSeid, s)
71   IN addrRegRange(use.vloc.vAddr, SEGsize(use.instance),
72     SEGinstanceInfo(use.instance).gl.growth);
73   $( gives the range of address registers used by a segment as a function of
74     the process id and the segment designator)
75
76 BOOLEAN noHole(seid pSeid; INTEGER size; virtualLocation vl; direction d;
77   SET OF segDes ssd) IS
78   NOT addrRegRange(vl.vAddr, size, d) SUBSET {0 .. MACmaxReg}
79   OR (EXISTS segDes s | s INSET ssd; useStruct use = SEGuseInfo(pSeid, s)
80     : use.vloc.idSpace = vl.idSpace
81     AND use.vloc.domain = vl.domain
82     AND addrRegRangeSeg(pSeid, s)
83     INTER addrRegRange(vl.vAddr, size, d) ~= {});
84   $(TRUE if a segment described by size, vl, and direction will NOT fit into
85     a hole in the address space designated by pSeid and ssd; this includes
86     testing for virtual memory underflow and overflow)
87
88   EXTERNALREFS
89
90   FROM mac:
91   INTEGER MACmaxVAddr, MACmaxOffset, MACmaxReg;
92
93   FROM smx:
94   seid: DESIGNATOR;
95   privType: {
96     privFileUpdateStatus,  privLink,
97     privLockSeg,           privModifyPriv,
98     privMount,             privSetLevel,
99     privStickySeg,         privSetPath,
100    privViolSimpSecurity,   privViolStarSecurity,
101    privViolSimpIntegrity,  privViolStarIntegrity,
102    privViolDiscrAccess,    privRealizeExecPermission,
103    privSignal,             privWalkPTable,
104    privHalt,               privKernelCall,
105    privViolCompartments,   privSetComm,
106    privImmigrate
107   };
108   daMode: {daRead, daWrite, daExecute};
109   securityCat: DESIGNATOR;
110   integrityCat: DESIGNATOR;
111   domainType: {nullDomain, kernelDomain, supervisorDomain, userDomain};
112   VFUN SENseidNsp(seid s) ^> INTEGER nsp;

```

```

113 VFUN TIIInfo(seid anySeid) ^> tIIStruct tIISt;
114 VFUN SMXflow(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
115 VFUN SMXdap(seid pSeid, oSeid; daType d) ^> BOOLEAN b;
116 VFUN SENSEidIndex(seid anySeid) ^> INTEGER index;
117
118 FROM pvm:
119 segDes: DESIGNATOR;
120 spaceType: {nullSpace, iSpace, dSpace};
121 direction: {up, down};
122 seid exampleSegmentSeid; $(used for segment creation)
123 INTEGER PVMmaxSegDes;
124 VFUN SEGInstanceInfo(seid segSeid) ^> instanceStruct is;
125 VFUN SEGUseInfo(seid pSeid; segDes segd) ^> useStruct us;
126 VFUN PVMsegmentSet(seid pSeid) ^> SET_OF segDes segSet;
127 VFUN PVMmappedSegmentSet(seid pSeid) ^> SET_OF segDes mappedSet;
128 VFUN SEGinUseIndexSet() ^> SET_OF INTEGER si;
129
130
131 ASSERTIONS
132
133 PVMimmVloc.domain = supervisorDomain;
134 PVMimmVloc.idSpace = iSpace;
135 PVMargVloc.domain = supervisorDomain;
136 PVMargVloc.idSpace = dSpace;
137 $( constraints on parameters)
138
139
140 FUNCTIONS
141
142 $(***** support for PSTreleaseProcess *****)
143
144 OFUN PVPreleaseProcessSupport(seid pSeid);
145 $( This function supports PRoreleaseProcess by deleting all segments in
146 the virtual memory named by "pSeid" and then deleting the virtual memory
147 itself)
148 DEFINITIONS
149 seid segSeid(segDes sd) IS SEGUseInfo(pSeid, sd).instance;
150 instanceStruct inst(segDes sd) IS SEGInstanceInfo(segSeid(sd));
151 ASSERTIONS
152 PVMvmExists(pSeid);
153 EFFECTS
154 FORALL segDes sd | SEGUseInfo(pSeid, sd) ~= ?
155 : 'SEGUseInfo(pSeid, sd) = ?
156 AND (IF inst(sd).refCount = 1 AND inst(sd).gl.sticky = FALSE
157 THEN 'SEGInstanceInfo(segSeid(sd)) = ?
158 AND 'TIIInfo(segSeid(sd)) = ?
159 ELSE
160 'SEGInstanceInfo(segSeid(sd))
161 = STRUCT(inst(sd).gl, inst(sd).growth,
162 inst(sd).refCount^1, inst(sd).data))
163 'PVMsegmentSet(pSeid) = ?;
164 'PVMmappedSegmentSet(pSeid) = ?;
165
166
167 $( NOTE ** There are two special segments that are appropriate to the invoke
168 and spawn operations: the argument segment "arg", already in the virtual

```

```

169 memory, which contains the data to be used by the initialized process;
170 and the immediate segment, "immSeid", which contains the code for the
171 process " this code is also called the process bootstrapper)
172
173 $(***** support for PSTinvoke *****)
174
175 OFUN PVPInvokeSupport(seid pSeid, immSeid; segDes arg); $(PVPInvokeSupport)
176 $( This function sets up the virtual memory of a process for
177 invocation; the new mapped set contains all previously mapped supervisor
178 segments and the argument and immediate segments; the argument segment
179 is mapped to a different virtual location, and a use is created for the
180 the immediate segment)
181 DEFINITIONS
182 instanceStruct immInst IS SEGInstanceInfo(immSeid);
183 useStruct argUse IS SEGUseInfo(pSeid, arg);
184 seid argSeid IS argUse.instance;
185 instanceStruct argInst IS SEGInstanceInfo(argSeid);
186 EXCEPTIONS
187 XbSgDes: argUse = ?;
188 XshrSg: argInst.gl.sharable;
189 XbSgSd: NOT (immInst ~= ?
190             AND SMXflow(pSeid, immSeid, {daRead}));
191 XsgNoAcc: NOT SMXdap(pSeid, immSeid, {daExecute});
192 XbSgRng:
193     NOT addrRegRange(PVMargVloc.vAddr, SEGsize(argSeid), argInst.gl.growth)
194     SUBSET {0 .. MACmaxReg};
195 XbSgRng:
196     NOT addrRegRange(PVMimmVloc.vAddr, SEGsize(immSeid), immInst.gl.growth)
197     SUBSET {0 .. MACmaxReg};
198 XpostEh:
199     nSegs(pSeid) >= PVMmaxSegDes;
200 ASSERTIONS
201 PVMvmExists(pSeid);
202 EFFECTS
203 'SEGUseInfo(pSeid, PVMargDes)
204   = STRUCT(argUse.instance, PVMargVloc, argUse.da);
205 $( create a reference to the immediate segment)
206 'SEGUseInfo(pSeid, PVMimmDes)
207   = STRUCT(immSeid, PVMimmVloc, {daRead, daExecute});
208 'SEGInstanceInfo(immSeid)
209   = STRUCT(immInst.gl, immInst.growth,
210           immInst.refCount + 1, immInst.data);
211 $( add the immediate segment to the address space)
212 'PVMsegmentSet(pSeid) = PVMsegmentSet(pSeid) UNION {PVMimmDes, PVMargDes};
213 $( unmap all segments except the argument segment,
214    and the immediate segment)
215 'PVMmappedSegmentSet(pSeid)
216   = {PVMargDes, PVMimmDes};
217 $( remap the argument segment)
218
219 $(***** support for PSTspawn *****)
220
221 OFUN PVPspawnSupport(seid parent, child; seid immSeid; segDes arg);
222                                     $(PVPspawnSupport)
223 $(creates a new address space named by child and inserts into it two
224 segment uses " the argument segment, arg, which is copied from the parent

```

```

225     address space, but occupies a different position " PVMargDes "
226     and the immediate segment, immSeid, which is shared)
227 DEFINITIONS
228     instanceStruct immInst IS SEGInstanceInfo(immSeid);
229     useStruct argUse IS SEGUseInfo(parent, arg);
230     seid argSeid IS argUse.instance;
231     instanceStruct argInst IS SEGInstanceInfo(argSeid);
232     seid argCopy
233         IS SOME seid s | SENseidNsp(s) = SENseidNsp(exampleSegmentSeid)
234             AND SEGInstanceInfo(s) = ?;
235     tiiStruct aTii IS TIIinfo(argSeid);
236     tiiStruct pTii IS TIIinfo(parent);
237     tiiStruct nTii IS STRUCT(aTii.nd, aTii.ms, pTii.owner, pTii.group,
238                             aTii.priv);
239 EXCEPTIONS
240     XbSgSd: NOT (immInst ~= ? AND SMXflow(parent, immSeid, {daRead}));
241     XsgNoAcc: NOT SMXdap(parent, immSeid, {daExecute});
242     XshrSg: argInst.gl.sharable;
243     XnWrtArgSg: NOT daWrite INSET argUse.da;
244     XbSgRng:
245         NOT addrRegRange(PVMargVloc.vAddr, SEGsize(argSeid), argInst.gl.growth)
246         SUBSET {0 .. MACmaxReg};
247     XbSgRng:
248         NOT addrRegRange(PVMimmVloc.vAddr, SEGsize(immSeid), immInst.gl.growth)
249         SUBSET {0 .. MACmaxReg};
250 RESOURCE_ERROR;
251 ASSERTIONS
252     PVMvmExists(parent);
253     NOT PVMvmExists(child);
254 EFFECTS
255     $( create a copy of the argument segment)
256     'TIIinfo(argCopy) = nTii;
257     'SEGInstanceInfo(argCopy) = STRUCT(argInst.gl, argInst.growth, 1,
258                                     argInst.data);
259     'SEGUseInfo(child, PVMargDes) = STRUCT(argCopy, argUse.vloc, argUse.da);
260     $( create a use for immSeid in child)
261     'SEGUseInfo(child, PVMimmDes)
262         = STRUCT(immSeid, PVMimmVloc, {daRead, daExecute});
263     'SEGInstanceInfo(immSeid)
264         = STRUCT(immInst.gl, immInst.growth, immInst.refCount + 1, immInst.data);
265     'PVMsegmentSet(child) = {PVMimmDes, PVMargDes};
266     'PVMmappedSegmentSet(child) = {PVMimmDes, PVMargDes};
267
268     $( "***** support for PROfork *****")
269
270 OFUN PVPforkSupport(seid parent, child);                                $(PVPforkSupport)
271     $(creates a new virtual memory, child, that is a copy of parent; some
272     segments are copied and others are merely shared; if a segment is
273     sharable in the parent process, it is not copied, but a use corresponding
274     to the instance in parent is created instead; if the segment is not
275     sharable, then a new instance of the segment is created, requiring the
276     allocation of an unused seid; in either case, corresponding segments have
277     identical segment designators in both processes; much mechanism in this
278     specification is devoted to describing the set of new seids created and
279     the mapping of this set onto the set of new segment instances)
280 DEFINITIONS

```



```

281     INTEGER nCopies
282     IS CARDINALITY
283       ({segDes sd | SEGinstanceInfo(SEGuseInfo(parent,
284                                     sd).instance).gl.sharable
285                                     = FALSE});
286     $(number of nonsharable segments in parent process)
287     SET OF INTEGER indexCopySet
288     IS SOME SET OF INTEGER si
289       | CARDINALITY(si) = nCopies AND si INTER SEGinUseIndexSet() = {};
290     SET OF seid copySet
291     IS {seid s | SENseidNsp(s) = SENseidNsp(exampleSegmentSeid)
292         AND SENseidIndex(s) INSET indexCopySet};
293     $(actual set of new seids)
294     useStruct use(segDes segd) IS SEGuseInfo(parent, segd);
295     instanceStruct inst(segDes segd)
296     IS SEGinstanceInfo(use(segd).instance);
297     EXCEPTIONS
298     RESOURCE_ERROR;
299     ASSERTIONS
300     PVMvmExists(parent);
301     NOT PVMvmExists(child);
302     EFFECTS
303     'PVMsegmentSet(child) = PVMsegmentSet(parent);
304     'PVMmappedSegmentSet(child) = PVMmappedSegmentSet(parent);
305     FORALL segDes segd | SEGuseInfo(parent, segd).instance ~= ?
306     : (IF inst(segd).gl.sharable
307       THEN
308         'SEGuseInfo(child, segd) = use(segd)
309         AND 'SEGinstanceInfo(use(segd).instance) =
310           STRUCT(inst(segd).gl, inst(segd).growth,
311                 inst(segd).refCount + 1, inst(segd).data)
312       ELSE
313         (LET seid copy INSET copySet
314          IN 'TIIinfo(copy) = TIIinfo(use(segd).instance)
315          AND 'SEGinstanceInfo(copy) =
316            STRUCT(inst(segd).gl, inst(segd).growth,
317                  1, inst(segd).data)
318          AND 'SEGuseInfo(child, segd) =
319            STRUCT(copy, use(segd).vloc, use(segd).da)
320          AND (FORALL segDes segd1 ~= segd
321              : 'SEGuseInfo(child, segd1).instance
322                ~= 'SEGuseInfo(child, segd).instance));
323         'SEGinUseIndexSet() = SEGinUseIndexSet() UNION indexCopySet;
324
325
326 END_MODULE

```

```

1  $(  
2      MODULE:      smx.specs (version 4.14)  
3      CONTENTS:    Security Model  
4      TYPE:        SPECIAL.specifications  
5      LAST CHANGED: 12/23/80, 17:40:14  
6  ")  
7  
8  MODULE smx  
9  
10  
11 $( This module now includes what used to be the contents of smx, prv, til,  
12    syl, and sen modules)  
13  
14     TYPES  
15  
16     $(from smx "" exportable)  
17 seid: DESIGNATOR;  
18 secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,  
19                    tExtent, tNull};  
20     privType: {  
21                 privFileUpdateStatus, privLink,  
22                 privLockSeg,          privModifyPriv,  
23                 privMount,            privSetLevel,  
24                 privStickySeg,        privSetPath,  
25                 privViolSimpSecurity, privViolStarSecurity,  
26                 privViolSimpIntegrity, privViolStarIntegrity,  
27                 privViolDiscrAccess,  privRealizeExecPermission,  
28                 privSignal,           privWalkPTable,  
29                 privHalt,             privKernelCall,  
30                 privViolCompartments, privSetComm,  
31                 privImmigrate  
32             };  
33  
34 daMode: {daRead, daWrite, daExecute};  
35 securityCat: DESIGNATOR;  
36 integrityCat: DESIGNATOR;  
37 domainType: {nullDomain, kernelDomain, supervisorDomain, userDomain};  
38  
39     $(from smx "" redeclarable)  
40 nonDisType: STRUCT_OF(  
41     INTEGER securityLevel; SET_OF securityCat securityCatS;  
42     INTEGER integrityLevel; SET_OF integrityCat integrityCatS);  
43     $(integrityCat is typically the null set)  
44 daType: SET_OF daMode;  
45 modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);  
46 tilStruct: STRUCT_OF(nonDisType nd;  
47     modeStruct ms; INTEGER owner, group; SET_OF privType priv);  
48  
49  
50     PARAMETERS  
51  
52 INTEGER SENmaxIndex $(maximum index component of a seid, 2^24 * 1),  
53     SENmaxNsp $(maximum nsp component of a seid, 2^8 * 1);  
54 INTEGER SENlowLevel, $(system low level)  
55     SENhighLevel; $(system high level)  
56

```

```

57
58     ASSERTIONS
59
60     FORALL seid s1; seid s2: (s1 = s2) = (SENseidNsp(s1) = SENseidNsp(s2)
61                                     AND SENseidIndex(s1) = SENseidIndex(s2));
62     $(this states that the nsp and index are isomorphic with the seid, and
63     thus uniquely identify it)
64
65     SENlowLevel < SENhighLevel;
66     $(defines the "greater than" relation for security in terms of the integer
67     relation ">")
68
69     SYLgetHigh() INSET {SENlowLevel .. SENhighLevel};
70     $(the current system high level is always within range)
71
72     FORALL seid s | TIIinfo(s) ~= ?
73     : LET nonDisType ndl = TIIinfo(s).nd
74       IN ndl.securityLevel INSET {SENlowLevel .. SENhighLevel}
75       AND ndl.integrityLevel INSET {SENlowLevel .. SENhighLevel};
76     $(restricts the values of the security and integrity levels for any
77     existing objects)
78
79
80     FUNCTIONS
81
82     $(***** sen " state functions *****)
83
84     VFUN SENseidNsp(seid anySeid) ^> INTEGER nsp;                $(SENseidNsp)
85     $( this is the nsp table entry component of a seid)
86     INITIALLY
87     nsp INSET {0 .. SENmaxNsp}; $(constrained by assertion above)
88
89     VFUN SENseidIndex(seid anySeid) ^> INTEGER index;            $(SENseidIndex)
90     $( this is the index component for a seid)
91     INITIALLY
92     index INSET {0 .. SENmaxIndex};
93     $(also characterized by assertion)
94
95     VFUN SENnspType(INTEGER nsp) ^> secureEntityType set;        $(SENnspType)
96     $( gives the type information as a function of the nsp component of
97     a seid)
98     INITIALLY NOT nsp INSET {0 .. SENmaxNsp} => set = ?;
99
100    $(***** sen " seid and nsp manipulation *****)
101
102    VFUN SENseidToInt(seid anySeid) ^> INTEGER i;                $(SENseidToInt)
103    $(gives the integer corresponding to a given seid)
104    DERIVATION
105    SENseidIndex(anySeid) + 2^24 * SENseidNsp(anySeid);
106
107    VFUN SENseidType(seid s) ^> secureEntityType set;            $(SENseidType)
108    $( returns the type information pertaining to a given seid)
109    DERIVATION
110    LET secureEntityType set1 = SENnspType(SENseidNsp(s))
111    IN IF set1 = ? THEN tNull ELSE set1;
112

```

```

113 VFUN SENmakeSeid(seid exampleSeid; INTEGER index) ^> seid rSeid;
114                                     $(SENmakeSeid)
115 $( forms a seid with an nsp the same as the example seid and the given
116   index; seid allocation is now done by the type managers for the
117   objects in question, allowing seids to be reused in the case of
118   objects that are dynamically allocated)
119 ASSERTIONS
120   index INSET {0 .. SENmaxIndex};
121 DERIVATION
122   SOME seid s | SENseidNsp(s) = SENseidNsp(exampleSeid)
123               AND SENseidIndex(s) = index;
124
125 $(***** syl functions *****)
126
127 VFUN SYLgetHigh() ^> INTEGER level;                                     $(SYLgetHigh)
128 $( represents the current highest security level for the system)
129 HIDDEN;
130 INITIALLY
131   level = SENhighLevel;
132
133 OFUN SYLsetHigh(INTEGER level);                                         $(SYLsetHigh)
134 $( sets the current highest security level for the system to the specified
135   value)
136 EXCEPTIONS
137   XbSysLev: NOT level INSET {SENlowLevel .. SENhighLevel};
138 EFFECTS
139   'SYLgetHigh() = level;
140
141 $(***** tii state function *****)
142
143 VFUN TIIinfo(seid s) ^> tiiStruct st;                                   $(TIIinfo)
144 $(returns the type-independent information for a system object; this
145   information includes discretionary, non-discretionary, and domain access
146   controls, privileges, and the owner and group for the object)
147 HIDDEN;
148 INITIALLY st = ?;
149
150 $(***** smx functions *****)
151
152 VFUN SMXhasPriv(seid pSeid; privType privl) ^> BOOLEAN b;             $(SMXhasPriv)
153 $( tells whether a given object " usually a process " has a particular
154   privilege)
155 DERIVATION
156   IF TIIinfo(pSeid) ~= ? THEN privl INSET TIIinfo(pSeid).priv ELSE FALSE;
157
158 VFUN SMXcompare(seid lSeid, hSeid) ^> BOOLEAN b;                       $(SMXcompare)
159 $( Primitive comparison telling whether lSeid <= hSeid. No privileges
160   are accounted for.)
161 DEFINITIONS
162   tiiStruct lTii IS TIIinfo(lSeid);
163   tiiStruct hTii IS TIIinfo(hSeid);
164 DERIVATION
165   IF lTii ~= ? AND hTii ~= ?
166   THEN   lTii.nd.securityLevel <= hTii.nd.securityLevel
167         AND lTii.nd.securityCatS SUBSET hTii.nd.securityCatS
168         AND lTii.nd.integrityLevel >= hTii.nd.integrityLevel

```

```

169         AND hTii.nd.integrityCatS SUBSET lTii.nd.integrityCatS
170     ELSE FALSE;
171
172 VFUN SMXcompWpriv(seid lSeid,hSeid; BOOLEAN priv1,priv2) *>BOOLEAN b;
173                                     $(SMXcompWpriv)
174     $(Primitive comparison telling whether lSeid <= hSeid, but
175     with privileges taken into account. See SMXflow comments re.
176     need for both SMXcompare and SMXcompWpriv.)
177 DEFINITIONS
178     tiiStruct lTii IS TIIinfo(lSeid);
179     tiiStruct hTii IS TIIinfo(hSeid);
180 DERIVATION
181     IF lTii ~= ? AND hTii ~= ?
182     THEN (priv1
183           OR
184           ( lTii.nd.securityLevel <= hTii.nd.securityLevel
185             AND lTii.nd.securityCatS SUBSET hTii.nd.securityCatS
186           )
187         )
188     AND
189     (priv2
190     OR
191     ( lTii.nd.integrityLevel >= hTii.nd.integrityLevel
192       AND hTii.nd.integrityCatS SUBSET lTii.nd.integrityCatS
193     )
194     )
195     ELSE FALSE;
196
197
198
199 VFUN SMXflow(seid pSeid, oSeid; daType d) *> BOOLEAN b;           $(SMXflow)
200     $( tells whether a given subject "pSeid" can access a given object "oSeid"
201     with the information flow specified by "flow", according to the
202     constraints of the military multilevel security model; these constraints
203     do not apply if the subject has the proper privilege)
204     $(" The reason for having both SMXcompare and SMXcompWpriv is that
205     the MLS tools have axiomatized SMXcompare, i.e., do not account
206     for privileges; the implementation does account for privileges.
207     The following relationships are true:
208     FORALL seid s1,s2(TIIinfo(s1)~=? AND TIIinfo(s2)~=? :
209         SMXcompWpriv(s1,s2,TRUE,TRUE) = TRUE
210         AND SMXcompWpriv (s1,s2,FALSE,FALSE) = SMXcompare(s1,s2)
211         AND SMXcompare(s1,s2) => SMXcompWpriv(s1,s2,TRUE,FALSE)
212         AND SMXcompare(s1,s2) => SMXcompWpriv(s1,s2,FALSE,TRUE);
213     ")
214 DERIVATION
215     (daWrite INSET i =>
216         SMXcompare(pSeid, oSeid)
217     OR
218         SMXcompWpriv(pSeid,oSeid,SMXhasPriv(pSeid,privViolStarSecurity),
219         SMXhasPriv(pSeid,privViolSimpIntegrity))
220     )
221 AND
222     (daRead INSET d =>
223         SMXcompare(oSeid, pSeid)
224     OR

```

```

225         SMXcompWpriv(oSeid,pSeid,SMXhasPriv(pSeid,privViolSimpIntegrity),
226             SMXhasPriv(pSeid,privViolStarIntegrity))
227     )
228     ;
229
230 VFUN SMXdap(seid pSeid, oSeid; daType d) ^> BOOLEAN b;           $(SMXdap)
231     $( tells whether a given subject "pSeid" can access a particular object
232       "oSeid" according to the discretionary access rules of the system ^^
233       similar to those of UNIX)
234     DEFINITIONS
235         tiiStruct p IS TIIinfo(pSeid);
236         modeStruct mst IS TIIinfo(oSeid).ms;
237         tiiStruct o IS TIIinfo(oSeid);
238         BOOLEAN access(daType requested, allowed)
239             IS requested
240                 SUBSET allowed
241                     UNION (IF SMXhasPriv(pSeid, privRealizeExecPermission)
242                         AND daExecute INSET allowed
243                         THEN {daRead}
244                         ELSE {});
245     DERIVATION
246         IF o ~= ?
247             THEN SMXhasPriv(pSeid, privViolDiscrAccess)
248                 OR access(d, mst.allMode)
249                 OR (p.group = o.group AND access(d, mst.groupMode))
250                 OR (p.owner = o.owner AND access(d, mst.ownerMode))
251             ELSE FALSE;
252
253     $(^^^^^^^^^^^^^^^^ tii ^^ extraction and insertion functions^^^^^^^^^^^^^^^^)
254
255 OFUN TIIcreateEntityLevel(seid oSeid; tiiStruct ntii);$(TIIcreateEntityLevel)
256     $( Initializes the tii information for an object "oSeid"; the object must
257       not have currently defined tii information; this is a service function
258       required for the creation of all types of objects in KSOS)
259     ASSERTIONS
260         TIIinfo(oSeid) = ?;
261         ntii.nd.securityLevel INSET {SENlowLevel .. SYLgetHigh()};
262         ntii.nd.integrityLevel INSET {SENlowLevel .. SENhighLevel};
263     EFFECTS
264         'TIIinfo(oSeid) = ntii;
265
266 VFUN TIIgetEntityLevel(seid pSeid, oSeid) ^> tiiStruct ntii;
267                                     $(TIIgetEntityLevel)
268     $(Retrieves the tii information of an object named by "oSeid", as
269       directed by process "pSeid"; mandatory and discretionary checks are also
270       performed; this function is used by
271       functions of the object maintaining modules, which provide status
272       information to the object that is concerned with getting and setting
273       object levels)
274     EXCEPTIONS
275         XnoObj: NOT(TIIinfo(oSeid) ~= ?
276                     AND SMXflow(pSeid, oSeid, {daRead}));
277     DERIVATION
278         TIIinfo(oSeid);
279
280 OFUN TIIsetEntityLevel(seid pSeid, oSeid; tiiStruct ntii);

```

```

281                                     $(TIIsetEntityLevel)
282 $( sets the typeaindependent information for an existing object
283 "oSeid" to the new value "ntii", as directed by process
284 "pSeid"; the privilege to set level is always required; if an
285 object's privileges are to be modified, the privilege to modify
286 privileges is required; because only privileged programs are allowed to
287 invoke this function, no other security checks a either mandatory
288 or discretionary a are made)
289 DEFINITIONS
290   tiiStruct otii IS TIIinfo(oSeid);
291   BOOLEAN sameOwner
292     IS TIIinfo(pSeid).owner = TIIinfo(oSeid).owner
293     OR SMXhasPriv(pSeid, privSetLevel);
294 EXCEPTIONS
295   $(privilege checking occurs in lev module)
296   XnSPriv:
297     otii.priv ~= ntii.priv AND NOT SMXhasPriv(pSeid, privModifyPriv);
298   XnoObj: otii = ?
299     OR (otii.ms ~= ntii.ms AND NOT SMXflow (pSeid,oSeid, {daWrite})
300        AND NOT sameOwner);
301   KEtiiNoSetLev:
302     otii.nd ~= ntii.nd AND NOT SMXhasPriv(pSeid,privSetLevel);
303   KEfiiNoSetPriv:
304     otii.priv ~= ntii.priv AND NOT SMXhasPriv(pSeid,privSetLevel);
305   KEtiiNoSetOwner:
306     otii.owner ~= ntii.owner AND NOT SMXhasPriv(pSeid,privSetLevel) AND (pSeid ~= oSeid,
307   KEtiiNoSetGroup:
308     otii.group ~= ntii.group AND NOT SMXhasPriv(pSeid,privSetLevel) AND (pSeid ~= oSeid,
309
310
311 ASSERTIONS
312   ntii.nd.securityLevel INSET {SENlowLevel .. SYLgetHigh()};
313   ntii.nd.integrityLevel INSET {SENlowLevel .. SENhighLevel};
314 EFFECTS
315   'TIIinfo(oSeid) = ntii;
316
317 OFUN TIIclearEntityLevel(seid oSeid);                                     $(TIIclearEntityLevel)
318 $( deletes the typeaindependent information for an object "oSeid"; this
319   function is used for implementing functions that delete objects a and
320   thus must also delete the tii info)
321 EFFECTS
322   'TIIinfo(oSeid) = ?;
323
324
325 END_MODULE

```

```

1  $("      MODULE:      spf.specs (version 4.6)
2          CONTENTS:      Special Functions
3          TYPE:          SPECIAL.specifications
4          LAST CHANGED:   7/28/80, 09:56:25
5  ")
6
7
8  MODULE spf
9
10
11 $( This module will contain the specifications for functions which do not
12    properly fit into any of the other kernel modules.  Functions so envisioned
13    will synchronize the file system, cause the system to halt, and set the
14    maximum system security level.  As of now these functions are left
15    unspecified)
16
17
18     TYPES
19
20 SPFFunctionType: {syncSPF, immSegSPF, sysHaltSPF, levelSetSPF};
21 SPFargs: VECTOR_OF INTEGER;
22
23
24     EXTERNALREFS
25
26     FROM smx:
27 seid: DESIGNATOR;
28
29
30     FUNCTIONS
31
32 OFUN SPFspecialFunction(seid pSeid; SPFFunctionType fn; SPFargs args);
33                                     $(SPFspecialFunction)
34     $(this function will be further specified when its exact nature is known)
35
36 END_MODULE

```


SECURE MINICOMPUTER OPERATING SYSTEM (KSOS) UNIX EMULATOR COMPUTER PROGRAM DEVELOPMENT SPECIFICATION (TYPE B5)

Department of Defense Kernelized Secure Operating System

Contract MDA 903-77-C-0333
CDRL 0002AG

Prepared for:

Defense Supply Service - Washington
Room 1D245, The Pentagon
Washington, DC 20310



**Ford Aerospace &
Communications Corporation**
Western Development
Laboratories Division

3939 Fabrian Way
Palo Alto, California 94303

NOTICE

The Department of Defense Kernelized Secure Operating System (KSOS) is being produced under contract to the U.S. Government. KSOS is intended to be compatible with the Western Electric Company's UNIX[®] Operating System (a proprietary product). KSOS is not part of the UNIX license software and use of KSOS is independent of any UNIX license agreement. Use of KSOS does not authorize use of UNIX in the absence of an appropriate licensing agreement.

This document is a major revision of an earlier draft dated July 1978. That and other previous editions are obsolete and should not be used.

UNIX and PWB/UNIX are trade/service marks of the Bell System.

DEC and PDP are registered trademarks of the Digital Equipment Corporation, Maynard MA.

TABLE OF CONTENTS

1.	SCOPE	1
1.1.	Identification	1
1.2.	Functional Summary	1
1.3.	Terminology	1
2.	APPLICABLE DOCUMENTS	5
2.1.	Government Documents	5
2.1.1.	Directives, Manuals and Standards	5
2.1.2.	Reports	5
2.2.	Non-Government Documents	6
2.3.	Other References Not Part of This Specification	7
3.	REQUIREMENTS	8
3.1.	Computer Program Definition	8
3.1.1.	Interface Requirements	8
3.1.1.1.	Interface Block Diagram	8
3.1.1.2.	Detailed Interface Definition	9
3.1.1.2.1.	Files	9
3.1.1.2.2.	Directories	10
3.1.1.2.3.	Processes	10
3.1.1.2.4.	Pipes	10
3.1.1.2.5.	Ports	10
3.2.	Detailed Functional Requirements	12
3.2.1.	abort	12
3.2.1.1.	Inputs	12
3.2.1.2.	Processing	12
3.2.1.3.	Outputs	12
3.2.2.	await	12
3.2.2.1.	Inputs	12
3.2.2.2.	Processing	13
3.2.2.3.	Outputs	13
3.2.3.	break	13
3.2.3.1.	Inputs	13
3.2.3.2.	Processing	13
3.2.3.3.	Outputs	14
3.2.4.	capac	14
3.2.4.1.	Inputs	14
3.2.4.2.	Processing	14
3.2.4.3.	Outputs	14
3.2.5.	chdir	14
3.2.5.1.	Inputs	14
3.2.5.2.	Processing	14
3.2.5.3.	Outputs	15
3.2.5.	chmod	15
3.2.6.1.	Inputs	15
3.2.6.2.	Processing	15
3.2.6.3.	Outputs	16

3.2.7.	chown	16
3.2.8.	close	16
3.2.8.1.	Inputs	16
3.2.8.2.	Processing	16
3.2.8.3.	Outputs	16
3.2.9.	creat	17
3.2.9.1.	Inputs	17
3.2.9.2.	Processing	17
3.2.9.3.	Outputs	17
3.2.10.	csw	18
3.2.11.	dup	18
3.2.11.1.	Inputs	18
3.2.11.2.	Processing	18
3.2.11.3.	Outputs	18
3.2.12.	eofp	19
3.2.12.1.	Inputs	19
3.2.12.2.	Processing	19
3.2.12.3.	Outputs	19
3.2.13.	exec	19
3.2.13.1.	Inputs	19
3.2.13.2.	Processing	20
3.2.13.3.	Outputs	20
3.2.14.	exit	20
3.2.14.1.	Inputs	20
3.2.14.2.	Processing	20
3.2.14.3.	Outputs	20
3.2.15.	fork	21
3.2.15.1.	Inputs	21
3.2.15.2.	Processing	21
3.2.15.3.	Outputs	21
3.2.16.	fstat	21
3.2.16.1.	Inputs	21
3.2.16.2.	Processing	21
3.2.16.3.	Outputs	22
3.2.17.	getal	22
3.2.17.1.	Inputs	22
3.2.17.2.	Processing	22
3.2.17.3.	Outputs	22
3.2.18.	getgid	23
3.2.18.1.	Inputs	23
3.2.18.2.	Processing	23
3.2.18.3.	Outputs	23
3.2.19.	getpid	23
3.2.19.1.	Inputs	23
3.2.19.2.	Processing	23
3.2.19.3.	Outputs	23
3.2.20.	getuid	24
3.2.20.1.	Inputs	24
3.2.20.2.	getuid	24
3.2.20.3.	Outputs	24
3.2.21.	gprocs	24
3.2.21.1.	Inputs	24
3.2.21.2.	Processing	24
3.2.21.3.	Outputs	25
3.2.22.	gtty	25
3.2.22.1.	Inputs	25

3.2.22.2.	Processing	25
3.2.22.3.	Outputs	25
3.2.23.	indir	26
3.2.23.1.	Inputs	26
3.2.23.2.	Processing	26
3.2.23.3.	Outputs	26
3.2.24.	kill	26
3.2.24.1.	Inputs	26
3.2.24.2.	Processing	26
3.2.24.3.	Outputs	26
3.2.25.	link	27
3.2.25.1.	Inputs	27
3.2.25.2.	Processing	27
3.2.25.3.	Outputs	27
3.2.26.	mknod	28
3.2.26.1.	Inputs	28
3.2.26.2.	Processing	28
3.2.26.3.	Outputs	28
3.2.27.	mount	28
3.2.28.	nice	28
3.2.28.1.	Inputs	28
3.2.28.2.	Processing	29
3.2.28.3.	Outputs	29
3.2.29.	open	29
3.2.29.1.	Inputs	29
3.2.29.2.	Processing	29
3.2.29.3.	Outputs	29
3.2.30.	pipe	30
3.2.30.1.	Inputs	30
3.2.30.2.	Processing	30
3.2.30.3.	Outputs	30
3.2.31.	port	30
3.2.31.1.	Inputs	30
3.2.31.2.	Processing	31
3.2.31.3.	Outputs	31
3.2.32.	profil	31
3.2.32.1.	Inputs	31
3.2.32.2.	Processing	32
3.2.32.3.	Outputs	32
3.2.33.	ptrace	32
3.2.33.1.	Inputs	32
3.2.33.2.	Processing	32
3.2.33.3.	Outputs	33
3.2.34.	read	33
3.2.34.1.	Inputs	33
3.2.34.2.	Processing	33
3.2.34.3.	Outputs	34
3.2.25.	seek	34
3.2.25.1.	Inputs	34
3.2.25.2.	Processing	34
3.2.25.3.	Outputs	34
3.2.36.	setgid	35
3.2.36.1.	Inputs	35
3.2.36.2.	Processing	35
3.2.36.3.	Outputs	35
3.2.37.	setuid	35

3.2.37.1.	Inputs	35
3.2.37.2.	Processing	35
3.2.37.3.	Outputs	36
3.2.38.	sfork	36
3.2.38.1.	Inputs	36
3.2.38.2.	Processing	36
3.2.38.3.	Outputs	36
3.2.39.	signal	36
3.2.39.1.	Inputs	36
3.2.39.2.	Processing	37
3.2.39.3.	Outputs	37
3.2.40.	sleep	38
3.2.40.1.	Inputs	38
3.2.40.2.	Processing	38
3.2.40.3.	Outputs	38
3.2.41.	stat	38
3.2.41.1.	Inputs	38
3.2.41.2.	Processing	38
3.2.41.3.	Outputs	39
3.2.42.	stime	39
3.2.43.	stty	39
3.2.43.1.	Inputs	39
3.2.43.2.	Processing	39
3.2.43.3.	Outputs	39
3.2.44.	sync	40
3.2.44.1.	Inputs	40
3.2.44.2.	Processing	40
3.2.44.3.	Outputs	40
3.2.45.	time	40
3.2.45.1.	Inputs	40
3.2.45.2.	Processing	40
3.2.45.3.	Outputs	40
3.2.46.	times	41
3.2.46.1.	Inputs	41
3.2.46.2.	Processing	41
3.2.46.3.	Outputs	41
3.2.47.	umount	41
3.2.48.	unlink	42
3.2.48.1.	Inputs	42
3.2.48.2.	Processing	42
3.2.48.3.	Outputs	42
3.2.49.	wait	42
3.2.49.1.	Inputs	42
3.2.49.2.	Processing	42
3.2.49.3.	Outputs	43
3.2.50.	write	43
3.2.50.1.	Inputs	43
3.2.50.2.	Processing	43
3.2.50.3.	Outputs	44
3.2.51.	Directory Manager	44
3.2.51.1.	Inputs	45
3.2.51.2.	Processing	45
3.2.51.3.	Outputs	46

3.2.52.	Network Interface	46
3.2.52.1.	Inputs	47
3.2.52.2.	Processing	48
3.2.52.2.1.	Initiation of Connections	48
3.2.52.2.2.	Flow Control	48
3.2.52.2.3.	Data Stream Integrity	48
3.2.52.2.4.	Urgent Processing	48
3.2.52.2.5.	Formatting	48
3.2.52.3.	Outputs	49
3.2.53.	Special Requirements	49
3.2.53.1.	Human Performance	49
3.2.53.2.	Government Furnished Property List	49
3.3.	Adaptation	49
4.	QUALITY ASSURANCE PROVISIONS	50
4.1.	Introduction	50
4.1.1.	Category I Testing	50
4.1.2.	Category II Testing	51
5.	PREPARATION FOR DELIVERY	52
6.	NOTES	53
7.	APPENDIX	54

**Kernelized Secure Operating System (KSOS):
UNIX Emulator Computer Program Development Specification (Type B5)**

1. SCOPE

1.1 Identification

This specification establishes the performance, design, development and test requirements for the Kernelized Secure Operating System (referred to as "KSOS") UNIX Emulator. The UNIX Emulator is intended to support the user program interface of the Western Electric UNIX operating system Version 6.0.

1.2 Functional Summary

The UNIX Emulator CPCI creates a system call interface which is compatible with that provided by the UNIX operating system [Ritchie 74], [Thompson 75]. The UNIX Emulator is an operating system, running on the "Kernel Machine", which creates UNIX level resources from the more primitive resources provided by the Kernel. The hierarchical UNIX file structure is created through the cooperation between the UNIX Emulator and the trusted Directory Manager. The Emulator also contains the per user aspects of support for the interface to a computer network.

The remainder of these specifications are organized as follows:

- ♦ Section 2 contains the references
- ♦ Section 3 contains the requirements for the CPCI
- ♦ Section 4 contains the quality assurance provisions for the CPCI
- ♦ Sections 5 and 6 are not applicable and are null.

These specifications serve two distinct purposes. The first is to define the UNIX Emulator in a generic, machine independent way. The second is to define an implementation for the PDP-11/70. Thus, there are many PDP-11/70 specific remarks. These should be treated as non-binding, informative commentary for non-PDP-11/70 implementations.

1.3 Terminology

The language used throughout this specification attempts to conform to the guidelines of Section 3.2.3 of MIL-STD-490. In particular, the word "shall" means that the specification expresses a provision that is binding. The words "should" and "may" mean that the specification expresses a provision which is non-mandatory. The word "will" is used to express a declaration of purpose on the part of the Government.

KSOS shall support the mandatory DoD security policy of DoD Directive 5200.1-R that is embodied in a Government approved mathematical model. [Verif

UNIX Emulator B5 Specification

78] Proofs of the system's security properties shall be in terms of this model. KSOS shall provide a minimum of eight (8) hierarchical security classifications categories (or simply security categories) and a minimum of thirty-two (32) mutually independent security compartments. The security categories shall be such that

UNCLASSIFIED < CONFIDENTIAL < SECRET < TOP SECRET

Where "<" is defined in accordance with the requirements of DoD 5200.1-R. One security compartment shall be reserved for read protection of system data bases and programs. This compartment shall be called the "system" compartment. KSOS shall provide for Kernel-enforced integrity. Integrity is defined as the formal mathematical dual of security [Biba 75]. At least four (4) hierarchical integrity classifications categories (or simply integrity categories) shall be provided in KSOS. The integrity categories include at least the following three classification categories:

USER < OPERATOR < ADMINISTRATOR

Although there is no immediate requirement for integrity compartments, the development contractor should include provisions to ease their later inclusion. For the remainder of this specification, the term "security level" means the combination of a security category and a set of security compartments (which may be null). The term "integrity level" means the combination of an integrity category, and a (presently always null) set of integrity compartments. The term "level" (or access level) means the security and integrity level, that is, the combination of a security category, a set of security compartments (which may be null), an integrity category, and a (presently always null) set of integrity compartments. The "access matrix" defines for a particular user or group the combination of the maximum security level permitted for that user or group, all security compartments to which that user or group has access, the maximum integrity level for that user or group, and all integrity compartments (presently always null) to which that user or group has access.

In a secure system it is necessary to have processes external to the kernel which perform security critical functions. These have come to be called "trusted processes". In KSOS there is a refinement in this terminology. As Figure 1-1 shows there are subdivisions of the software inside the security perimeter. Some of the processes are "privileged". This means that they may violate one or more of the security rules enforced by the Kernel. These processes must be designed and implemented with the same care and thoroughness as the Kernel itself. A second subdivision has been termed "responsible" processes. These are processes are not privileged to violate any of the Kernel's rules, but which are nonetheless important to the overall security of the system. Examples of responsible processes include the Secure Mail facility and the processes which manipulate the data bases describing the users' security and integrity permissions (the "access matrix"). These processes must also be carefully designed and implemented, but do not require formal specification of verification. Because they do perform security-related functions, they cannot include untrusted components (e.g. the UNIX Emulator). The third subdivision is the "encapsulated utilities" (EU in the figure). These are self contained

UNIX Emulator B5 Specification




Figure 1-1. Terminology Figure

functions to which the user simply supplies parameters. The encapsulated utilities are not privileged to violate Kernel rules, and do not affect the security of users. The chief function in the encapsulated utilities is system

UNIX Emulator B5 Specification

generation.

UNIX Emulator B5 Specification

2. APPLICABLE DOCUMENTS

The following documents, of exact issue shown, form a part of this specification to the extent specified herein. In the event of a conflict between the referenced documents and the contents of this specification, this specification shall be considered a superseding requirement. In the text references to these documents are in the form [Name date], e.g. [Biba 75].

2.1 Government Documents

2.1.1 Directives, Manuals and Standards

- a. DoD 5200.1-R Information Security Program Regulation
- b. DoD 5200.28 Security Requirements for Data Processing (ADP)
- c. DoD 5200.28-M ADP Security Manual
- d. MIL-STD-483 Configuration Management
- e. MIL-STD-490 Specification Practices
- f. MIL-STD-1521A Technical Reviews and Audits

2.1.2 Reports

- a. [A-Specs 78] "KSOS System Specification (Type A)", WDL-TR7808 Revision 1, Ford Aerospace and Communications Corporation, Palo Alto, CA (July 1978).
- b. [Bell and LaPadula 73] Bell, D.E. and LaPadula, L.J., "Secure Computer Systems", ESD-TR-73-278, Volume I-III, MITRE Corporation, Bedford, MA (November 1973 - June 1974).
- c. [Biba 75] Biba, K.J., "Integrity Considerations for Secure Computer Systems", MTR-3153, MITRE Corporation, Bedford, MA (June 1975).
- d. [Kernel 78] "KSOS Computer Program Development Specification (Type B5): Security Kernel", WDL-TR7932, Ford Aerospace and Communications Corporation, Palo Alto, CA (September 1978).
- e. [NKSr 78] "KSOS Computer Program Development Specification (Type B5): Non-Kernel Security-Related Software", WDL-TR7934, Ford Aerospace and Communications Corporation, Palo Alto, CA (September 1978).
- f. [Verif 78] "KSOS Verification Plan", WDL-TR7809, Ford Aerospace and Communications Corporation, Palo Alto, CA (March 1978).
- g. [Walter et al. 74] Walter, K.G. et al., "Primitive Models for Computer Security", ESD-TR-74-117, Case Western Reserve University, Cleveland, OH (January 1974).

UNIX Emulator B5 Specification

2.2 Non-Government Documents

- a. [BBN 75] "Interface Message Processor, Specifications for the Interconnection of a Host and an IMP", Report 1822, Bolt Beranek and Newman Inc., Cambridge, MA, (December 1975).
- b. [Holmgren et al. 77] Holmgren, S.F., Healy, D.C., Jones, P.B. and Kasprzycki, E., "Illinois Inter-Process Communication Facility for UNIX", CAC Technical Memorandum 84, CCTC-WAD Document 7507, Center for Advanced Computation, University of Illinois at Urbana-Champaign, Urbana, IL (April 1977).
- c. [Lampson 73] Lampson, B., "A Note on the Confinement Problem", CACM, Volume 16, Number 10, pp 613 - 615 (October 73).
- d. [Lipner 75] Lipner, S.B., "A Comment on the Confinement Problem", Proc. Fifth Symposium on Operating Systems Principles, ACM SIGOPS Review, Volume 9, Number 5, pp 192 - 196 (19-21 November 1975).
- e. [Nemeth et al. 77] Nemeth, A.G., Sunshine, C., Zucker, S. and Tepper, S., "Progress Towards a DeFacto DoD UNIX Inter-Process Communication Standard", Working Paper, (May 1977). (Available from the authors.)
- f. [Parnas 72] Parnas, D.L., "A Technique for Software Module Specification with Examples", CACM, Volume 15, Number 5, pp 330 - 336 (May 1972).
- g. [Postel 80a] Postel, J.B., "DOD Standard Internet Protocol", (RFC 760, IEN 128), Information Sciences Institute, University of Southern California, Marina del Rey, CA (January 1980).
- h. [Postel 80b] Postel, J.B., "DOD Standard Transmission Protocol", (RFC 761, IEN 129), Information Sciences Institute, University of Southern California, Marina del Rey, CA (January 1980).
- i. [Ritchie 74] Ritchie, D.M. and Thompson, K., "The UNIX Timesharing System", CACM, Volume 17, Number 5, pp 365 - 375 (May 1974).
- j. [Roubine and Robinson 77] Roubine, O., L. Robinson, Special Reference Manual, 3rd ed., Technical Report CSG-45, SRI International, Menlo Park, CA (January 1977).
- k. [Thompson 75] Thompson, K. and Ritchie, D.M., "UNIX Programmer's Manual", Sixth Edition, Western Electric Corporation, Greensboro, NC (May 1975).
- l. [Sunshine 77] Sunshine, C., "Interprocess Communication Extensions for the UNIX Operating System: I Design Considerations", R-2064/1-AF, RAND Corporation,
- m. [Zucker 77] Zucker, S., "Interprocess Communication Extensions for the UNIX Operating System: II Implementation", R-2064/2-AF, RAND Corporation, Santa Monica, CA (June 1977).

UNIX Emulator B5 Specification

2.3 Other References Not Part of This Specification

- a. [UNIX 75] "UNIX Program Listings", Version 6.0, Western Electric Corporation, Greensboro, NC (May 1975).

3. REQUIREMENTS

3.1 Computer Program Definition

This CPCI defines an operating system program to be implemented upon the KSOS Security Kernel which shall emulate the call interface of the Western Electric Corp. UNIX operating system. The resulting operating system shall provide a UNIX environment which enforces DoD military security policy. For additional product specification information, the reader is referred to the KSOS System Specification (Type A) [A-Specs 78].

The KSOS UNIX Emulator has two major components: the operating system call Emulator and the Directory Manager.

3.1.1 Interface Requirements

The UNIX Emulator shall execute in the memory and privilege domain between user programs and the Security Kernel. On the PDP-11/70, the Emulator shall execute in the Supervisor domain of the machine. The Emulator data base shall be protected from tampering by the user process through the virtual memory hardware of the host computer. To the Security Kernel, the Emulator and the user program are considered to be one, untrusted process that simultaneously occupies more than one virtual memory domain.

The user program shall be capable of trapping to either the Emulator or directly to the Security Kernel. In the PDP-11/70 implementation of KSOS, the EMT instruction shall be used to trap to the Kernel. The TRAP instruction shall cause a trap to the Emulator. The Emulator shall be dispatched to the appropriate trap handler, saving the state of the user's registers. The PDP-11/70 KSOS system may allocate general register set 1 to the Emulator and user programs. General register set 0 may be used by the Security Kernel.

The relationship between the user program and the Emulator may be a cooperative one. Because the user may execute Kernel calls directly, it may be possible to circumvent the Emulator and cause undesirable effects. These effects shall be confined to the user process only without disturbing the processes of other system users except as other user processes have arranged and expect to cooperate with the corrupted process (for example through ports or files) and to the extent that the corrupted process deviates from its expected behavior. Although the Emulator data bases shall be protected from direct manipulation by the user program, the Emulator may be viewed as an extension of the user process, acting on the behalf of the user. The user may abuse the services of the Emulator at the peril of the process. Information compromises shall never result from erroneous interaction of the user process part and the Emulator.

3.1.1.1 Interface Block Diagram

Figure 1-2 is the interface block diagram for the KSOS UNIX Emulator. This diagram illustrates the layers of virtual machine or functional capability in the KSOS operating system. The KSOS UNIX Emulator interfaces directly with the Security Kernel and user programs through the interdomain call facility of the PDP-11/70 hardware. The PDP-11 EMT instruction shall be used for direct

UNIX Emulator B5 Specification




Figure 1-2. Interface Diagram

Emulator to Kernel communication. The PDP-11 TRAP and IOT instructions shall be used to emulate the UNIX call interface. Existing UNIX software uses the TRAP and IOT instructions for operating system calls. Hence, compatibility with current UNIX shall be maintained. The UNIX call Emulator spawns a directory manager as needed. Directory creation, deletion and update requests shall be transferred through a parameter segment when the Directory Manager is spawned by the Emulator. Responses from the Directory Manager (i.e., error reports) shall be sent by IPC message to the Emulator of the process which initiated the directory operation.

3.1.1.2 Detailed Interface Definition

3.1.1.2.1 Files

UNIX files shall remain intact in KSOS. The user shall be restricted from observing certain file status information, i.e., link counts. These restrictions shall be required to remove clandestine information channels between users of different security and integrity levels.

The Security Kernel shall support a mutual exclusion mechanism on Kernel level files. (This mechanism may be used to regulate updates on directories.) Future implementations of the KSOS Emulator should be able to take advantage of the open for exclusive use feature of the Kernel. When a file is opened for exclusive use, other processes attempting to open the file for reading or writing shall be forced to block until the file has been closed or until a system default timeout has expired. Support of exclusive use files at the KSOS user level is beyond the scope of the current Emulator contract, but would be a rewarding future addition to the Emulator.

UNIX Emulator B5 Specification

3.1.1.2.2 Directories

UNIX directory semantics shall be retained in their entirety in KSOS. The operation of the UNIX calls link, unlink and creat shall be compatible with standard UNIX relative to directory semantics. (However, the access checking on link and unlink shall be tightened in KSOS.)

The integrity of the KSOS file systems should be improved over the present Version 6.0 UNIX. Directories may be implemented using the Kernel file subtype mechanism, preventing casual manipulation of file directories. The Directory Manager shall provide a substantial demonstration of KSOS file subtypes. The KSOS user shall be unaware of the existence of subtypes; the Emulator shall make the directory mechanism transparent to the user process.

3.1.1.2.3 Processes

Standard UNIX process semantics shall also be retained. Due to the mediation of the Kernel, the invocation of trusted software and core dumping shall be secure. Trusted software shall be invoked through the Kernel K_invoke and K_spawn calls. Because core dumps shall be performed outside the security perimeter of the Kernel by the Emulator, unauthorized disclosure of information shall be prevented. The Emulator shall not write a core dump file without the appropriate access capability. Communication with trusted children shall be disallowed by the child, preventing compromises through ptrace and other inter-process communication (IPC) channels.

KSOS users may be capable of sharing data through common process segments. Process segments shall be manipulated through the direct user to Kernel interface. Future enhancements to the Emulator should provide Emulator calls for manipulating segments, preventing Emulator "blind spots" and providing a more congenial interface to the user for segment manipulation. Although the user may succeed in confusing the Emulator by changing process segments directly, no compromise of information shall occur.

3.1.1.2.4 Pipes

Version 6.0 pipes shall be implemented by the KSOS Emulator. Pipes may be implemented using Kernel-level shared segments.

3.1.1.2.5 Ports

Ports will be a user level mechanism for inter-process communication. KSOS ports will be based upon RAND ports, augmented by two new operating system calls (await and capac) which provide low level channel synchronization [Nemeth et al. 77]. Ports may be implemented using shared segment communications, providing greater channel bandwidth than current implementations. Further information on RAND ports may be found in [Sunshine 77] and [Zucker 77].

UNIX Emulator B5 Specification

abort	3.2.1.	IOT
await	3.2.2.	59
break	3.2.3.	17
capac	3.2.4.	63
chdir	3.2.5.	12
chmod	3.2.6.	15
chown	3.2.7.	(Subsumed into NKSR software)
close	3.2.8.	6
creat	3.2.9.	8
csw	3.2.10.	(unsupported)
dup	3.2.11.	38
eofp	3.2.12.	60
exec	3.2.13.	11
exit	3.2.14.	1
fork	3.2.15.	2
fstat	3.2.16.	28
getal	3.2.17.	58
getgid	3.2.18.	47
getpid	3.2.19.	20
getuid	3.2.20.	24
gprocs	3.2.21.	61
gtty	3.2.22.	32
indir	3.2.23.	0
kill	3.2.24.	37
link	3.2.25.	9
mknod	3.2.26.	14
mount	3.2.27.	(Subsumed into NKSR software)
nice	3.2.28.	34
open	3.2.29.	5
pipe	3.2.30.	42
port	3.2.31.	57
profil	3.2.32.	44
ptrace	3.2.33.	26
read	3.2.34.	3
seek	3.2.35.	19
setgid	3.2.36.	46
setuid	3.2.37.	23
sfork	3.2.38.	62
signal	3.2.39.	48
sleep	3.2.40.	35
stat	3.2.41.	18
stime	3.2.42.	(Subsumed into NKSR Software)
stty	3.2.43.	31
sync	3.2.44.	36
time	3.2.45.	13
times	3.2.46.	43
umount	3.2.47.	(Subsumed into NKSR Software)
unlink	3.2.48.	10
wait	3.2.49.	7
write	3.2.50.	4

Figure 3-3. List of KSOS UNIX Emulator Calls

3.2 Detailed Functional Requirements

This paragraph describes the detailed functional requirements of the KSOS UNIX Emulator. Figure 3-2 summarizes the Emulator calls by name (listed alphabetically), Specification paragraph number, and system entry number. The system entry number corresponds to the low order byte of the PDP-11/70 TRAP instruction used to enter the Emulator.

For each call, a list of exception conditions is provided. The order of this list is not necessarily the order in which the conditions would be evaluated. The development contractor may alter the order or add to the exception conditions as needed, subject to Government approval. Various exceptions detected during a system call may not return distinct error numbers, but rather, may maintain compatibility with existing UNIX error numbers.

3.2.1 abort

3.2.1.1 Inputs

The Emulator call abort shall take the following parameters.

None

3.2.1.2 Processing

The Emulator abort call shall cause a core image of the user-mode part of the process and the system's per-user data for that user to be dumped to a disk file. The core image file shall not be produced unless all the conditions necessary to create a file in the working directory of the process have been satisfied, including mandatory security and discretionary access permissions. (These same conditions shall apply to dumps created as a result of some process induced fault.)

3.2.1.3 Outputs

The Emulator call abort shall return the following values after execution.

None (Does not return)

The core file shall be created, if possible, and the parent process shall be notified about the death of this process.

3.2.2 await

3.2.2.1 Inputs

The Emulator call await shall take the following parameters.

to: time out
wc: wake up conditions

UNIX Emulator B5 Specification

3.2.2.2 Processing

(Not present in Version 6.0 UNIX.) The Emulator call `await` shall permit a process to suspend its execution until either a time out event has occurred or one of the specified wake up conditions has been met. The time out value may be specified in 1/60 seconds, or some other appropriate time unit. The minimum set of wake up conditions to be implemented shall be:

- a. wake up when any pipe is written
- b. wake up when any pipe is read
- c. wake up when any pipe becomes empty

If the desired wake up event has occurred at the time of the `await` call, the process shall not block, but return to the process immediately. The wake up value zero shall cause the process to wait only on the wake up conditions.

3.2.2.3 Outputs

The Emulator call `await` shall return the following values after execution.

cw: wake up event which caused process to awaken
ec: error conditions

The Emulator call `await` shall detect and report the following exception (error) conditions.

EX: invalid time out value
EX: invalid wake up condition specified

3.2.3 break

3.2.3.1 Inputs

The Emulator call `break` shall take the following parameters.

na: new break address

3.2.3.2 Processing

The Emulator call `break` shall increase the length of the process data segment to the new break address. Break shall also permit process data segments to shrink. Access attempts to locations between the break address and the stack pointer shall cause a segmentation violation signal if access is attempted to a currently unmapped portion of the virtual address space (PDP-11/70 peculiarity). The PDP-11/70 implementation of KSOS shall express the break address as a multiple of 512 bytes.

3.2.3.3 Outputs

The Emulator call break shall return the following values after execution.

ec: error conditions

The Emulator call break shall detect and report the following exception (error) conditions.

EX: physical memory resources exhausted

EX: memory management hardware resources exhausted

3.2.4 capac

3.2.4.1 Inputs

The Emulator call capac shall take the following parameters.

fd: file descriptor of pipe (or port)

cv: pointer to two integer vector for byte counts

3.2.4.2 Processing

(Not present in Version 6.0 UNIX.) (This call subsumes the RAND empty call [Sunshine 77] and [Zucker 77].) The capac Emulator call shall return the number of bytes in a pipe (or port) available for reading and writing. This call shall not cause the process to block.

3.2.4.3 Outputs

The Emulator call capac shall return the following values after execution.

br: number of bytes available for reading

bw: number of bytes available for writing

The Emulator call capac shall detect and report the following exception (error) conditions.

EX: invalid file descriptor

EX: file descriptor does not refer to pipe (or port)

EX: improper vector pointer address

3.2.5 chdir

3.2.5.1 Inputs

The Emulator call chdir shall take the following parameters.

UNIX Emulator B5 Specification

pn: pathname of directory file

3.2.5.2 Processing

The Emulator call chdir shall change the current process working directory to the specified directory.

3.2.5.3 Outputs

The Emulator call chdir shall return the following values after execution.

ec: error conditions

The Emulator call chdir shall detect and report the following exception (error) conditions.

EX: directory required to interpret pathname is unmounted
EX: file on name interpretation path is not a directory
EX: directory on name interpretation path does not exist
EX: directory on name interpretation path has search permission denied
EX: directory required for name interpretation has higher security level
 (returned as if directory did not exist)
EX: directory on name interpretation path is at lower integrity level
 (returned as if directory did not exist)

3.2.6 chmod

3.2.6.1 Inputs

The Emulator call chmod shall take the following parameters.

pn: pointer to file pathname
da: new discretionary access mode

3.2.6.2 Processing

The chmod call shall change the discretionary access mode of a file. Only the file owner and sufficiently privileged processes may change the discretionary access mode of a file. Because the chmod call is partially processed by the untrusted Emulator, it is conceivable, though unlikely, that there would be a "trojan horse" embedded in the Emulator that would not actually perform the requested chmod. For example, the trojan horse could always set the protection to 777. If the user wishes to be absolutely sure that the chmod is done only by trusted software, the File Access Modifier [NKSr 78] may be employed at command level. These remarks apply only to the discretionary access information. The hypothetical trojan horse could not affect the security level of the file.

UNIX Emulator B5 Specification

3.2.6.3 Outputs

The Emulator call chmod shall return the following values after execution.

ec: error conditions

The Emulator call chmod shall detect and report the following exception (error) conditions.

EX: directory needed for name interpretation is on unmounted file system
EX: file on name interpretation path is not a directory
EX: directory on name interpretation path does not exist
EX: directory on name interpretation path has search permission denied
EX: directory on interpretation path is at higher security level (returned as if directory did not exist)
EX: directory on interpretation path is at lower integrity level (returned as if directory did not exist)
EX: file name could not be found in directory
EX: file is at lower security level than process (returned as if file did not exist)
EX: file is at higher integrity level than calling process (returned as if file did not exist)
EX: file does not exist

3.2.7 chown

The standard UNIX call chown has been subsumed into the File Access Modification process of the Non-Kernel Security Related Software CPCI.

3.2.8 close

3.2.8.1 Inputs

The Emulator call close shall take the following parameters.

fd: file descriptor

3.2.8.2 Processing

The close call shall terminate input/output to the specified file, device, or pipe. A file descriptor shall specify the file to be closed. In the PDP-11/70 implementation, the file descriptor shall be passed in register R0. The last close of a file whose link count has decremented to zero shall cause the file to be deleted. Termination of a process shall implicitly close all of its open files.

3.2.8.3 Outputs

The Emulator call close shall return the following values after execution.

UNIX Emulator B5 Specification

ec: error conditions

The Emulator call close shall detect and report the following exception (error) conditions.

EX: invalid file descriptor

3.2.9 creat

3.2.9.1 Inputs

The Emulator call creat shall take the following parameters.

pn: pointer to user pathname string in user domain
da: discretionary access mode of new file

3.2.9.2 Processing

The Emulator call creat shall create a new file. The file shall be subsequently accessible by

- a. the completely specified (rooted) pathname, or
- b. a concatenation of the process working directory string and the specified pathname string.

The file shall be created with the specified discretionary access. (See the remark in chmod about discretionary access.) The parent directory of the file shall be updated to include the new file. The newly created file shall have a security and integrity level equal to that of the calling process.

If the file exists, but is at a security level where it cannot be written, the error return shall be as if discretionary access to it was denied. This preserves the UNIX construct of lock files.

An open file descriptor shall be returned by creat. This descriptor shall be used to reference the file in subsequent operations. The file shall be opened for writing. In the PDP-11/70 implementation, the file descriptor shall be returned in register R0.

3.2.9.3 Outputs

The Emulator call creat shall return the following values after execution.

fd: open file descriptor
ec: error conditions

UNIX Emulator B5 Specification

The Emulator call `creat` shall detect and report the following exception (error) conditions.

- EX: security or integrity level not permitted in file system upon which the file would be created
- EX: file system required for pathname interpretation is not mounted
- EX: file name required for pathname interpretation is not a directory
- EX: directory in name interpretation path is unsearchable or does not exist
- EX: file to be created exists and is a directory
- EX: directory to be written is not writably mounted
- EX: file does not exist and write permission on directory is denied
- EX: file does exist and write permission on it is denied
- EX: too many open files

3.2.10 csw

The standard UNIX operating system call `csw` will not be supported by the KSOS Kernel.

3.2.11 dup

3.2.11.1 Inputs

The Emulator call `dup` shall take the following parameters.

fd: file descriptor

3.2.11.2 Processing

The `dup` call shall allocate and return another file descriptor which is synonymous with the original descriptor. The assignment algorithm shall return the lowest value available.

3.2.11.3 Outputs

The Emulator call `dup` shall return the following values after execution.

fd: new file descriptor
ec: error conditions

The Emulator call `dup` shall detect and report the following exception (error) conditions.

- EX: invalid file descriptor
- EX: no free file descriptors

3.2.12 eofp

3.2.12.1 Inputs

The Emulator call eofp shall take the following parameters.

fd: file descriptor of pipe

3.2.12.2 Processing

(Not present in Version 6.0 UNIX.) The Emulator call eofp shall write an end of file on a pipe. Additional writes may be performed on the pipe, but they are suspended until the eof has been read. If the eofp call has been issued, after reading all the valid data, the next read on the pipe shall return 0 characters.

3.2.12.3 Outputs

The Emulator call eofp shall return the following values after execution.

ec: error conditions

The Emulator call eofp shall detect and report the following exception (error) conditions.

EX: invalid file descriptor or file descriptor does not refer to a pipe

3.2.13 exec

3.2.13.1 Inputs

The Emulator call exec shall take the following parameters.

na: pointer to file name string of process prototype image

ap: pointer to zero terminated list of argument strings

3.2.13.2 Processing

The Emulator call exec shall overlay the current user domain process image with a new image as defined by the process prototype file. The calling process image shall be lost. Open files shall remain open after the execution of the exec call. Signal function values shall be set to the defaults, except for ignored signals, which remain ignored. Profiling shall be disabled by the execution of exec. If the process image to be incarnated is privileged (i.e., privilege bits in the type independent information for the file are enabled) the effective process user ID shall be updated accordingly. Reference [Thompson 75] section A.OUT(V) for the definition of the process prototype file format.

UNIX Emulator B5 Specification

3.2.13.3 Outputs

The Emulator call `exec` shall return the following values after execution.

ec: error conditions

The Emulator call `exec` shall detect and report the following exception (error) conditions.

- EX: security level of process is too low to execute (read) file (returned as if file did not exist)
- EX: integrity level of process is too high to execute (read) file (returned as if file did not exist)
- EX: process prototype file does not exist
- EX: process user ID does not have execute permission to prototype file
- EX: too many arguments (argument list too large) The limit on the size of the argument segment shall be no smaller than the present limit (approximately 512 characters).
- EX: file prototype format error
- EX: not enough memory resources to execute process image

3.2.14 exit

3.2.14.1 Inputs

The Emulator call `exit` shall take the following parameters.

es: exit status value

3.2.14.2 Processing

The Emulator call `exit` shall cause the process (both user and Emulator parts) to terminate execution. All open files shall be closed. The parent process shall be notified of the process termination. The exit status value shall be returned to the parent as part of the notification. The `exit` call shall never return.

3.2.14.3 Outputs

The Emulator call `exit` shall return the following values after execution.

None

The Emulator call `exit` shall detect and report the following exception (error) conditions.

None

UNIX Emulator B5 Specification

3.2.15 fork

3.2.15.1 Inputs

The Emulator call fork shall take the following parameters.

None

3.2.15.2 Processing

The Emulator call fork shall create a new process. The new process image shall be identical to the parent (original) process image. Open files shall be passed to the child (newly created) process. The child process shall receive the signal function values of the parent process. If profiling is enabled in the parent process, profiling shall be inherited by the child. The return value of the parent process shall be the process ID of the child process. The child process shall receive the return value zero. In the PDP-11/70 implementation of KSOS, the program counter of the parent process shall be incremented by two, causing the parent to return to a different location after the execution of fork. The clandestine information channel offered by the resource exhaustion condition shall be bandwidth limited by the KSOS Security Kernel.

3.2.15.3 Outputs

The Emulator call fork shall return the following values after execution.

pi: process id of child (returned to parent process), zero (returned to child process)

The Emulator call fork shall detect and report the following exception (error) conditions.

EX: process space exhausted

3.2.16 fstat

3.2.16.1 Inputs

The Emulator call fstat shall take the following parameters.

fd: file descriptor (passed in register R0)

ba: buffer address for file status information

3.2.16.2 Processing

The fstat call shall return the same information (except for shared values like link counts that could be used as channels) and in the same format as the stat call below. The file descriptor shall refer to a presently open (read or read/write mode) file. In KSOS, a file must be readable with respect to

UNIX Emulator B5 Specification

mandatory security before the control information can be accessed to prevent clandestine information channels.

3.2.16.3 Outputs

The Emulator call fstat shall return the following values after execution.

fs: file status information
ec: error conditions

The Emulator call fstat shall detect and report the following exception (error) conditions.

EX: invalid file descriptor
EX: file has not been opened for reading

3.2.17 getal

3.2.17.1 Inputs

The Emulator call getal shall take the following parameters.

bp: pointer to buffer for access level information

3.2.17.2 Processing

(Not present in Version 6.0 UNIX.) The getal call shall return the type independent information for the process: security classification, security category, integrity level, integrity category, and effective/real group and user identifiers. The user identification information is returned in the full bit-precision supported by the Security Kernel. The user identification returned by the getuid and getgid calls represents the shorter "mapped" identifiers as transformed by the Emulator.

3.2.17.3 Outputs

The Emulator call getal shall return the following values after execution.

al: process access level information

The Emulator call getal shall detect and report the following exception (error) conditions.

EX: invalid buffer pointer address

UNIX Emulator B5 Specification

3.2.18 getgid

3.2.18.1 Inputs

The Emulator call getgid shall take the following parameters.

None

3.2.18.2 Processing

The Emulator call getgid shall return the effective and real group ID of the process. In the PDP-11/70 implementation of KSOS, the high order byte of R0 shall contain the effective ID. The low order byte shall contain the real ID. The real ID shall identify the user who initiated execution of the process.

The KSOS Security Kernel supports user and group ID's which are longer than the 8-bits implemented in standard UNIX. The UNIX Emulator shall map longer Kernel names to shorter UNIX compatible identifiers. (This note shall also apply to the Emulator call getuid, below.)

3.2.18.3 Outputs

The Emulator call getgid shall return the following values after execution.

er: effective and real group ID of process

The Emulator call getgid shall detect and report the following exception (error) conditions.

None

3.2.19 getpid

3.2.19.1 Inputs

The Emulator call getpid shall take the following parameters.

None

3.2.19.2 Processing

The Emulator call getpid shall return the process ID of the calling process. This name may be a mapped version of the longer Kernel names for processes (SEID's).

3.2.19.3 Outputs

The Emulator call getpid shall return the following values after execution.

UNIX Emulator B5 Specification

pi: process ID of calling process

The Emulator call getpid shall detect and report the following exception (error) conditions.

None

3.2.20 getuid

3.2.20.1 Inputs

The Emulator call getuid shall take the following parameters.

None

3.2.20.2 getuid

The Emulator call getuid shall return the effective and real user ID of the calling process. In the PDP-11/70 implementation of KSOS, the low order byte of R0 shall contain the real user ID. The high order byte shall return the effective user ID. See the notes in getpid above for a discussion of the uses of real and effective ID.

3.2.20.3 Outputs

The Emulator call getuid shall return the following values after execution.

er: effective and real user ID of calling process

The Emulator call getuid shall detect and report the following exception (error) conditions.

None

3.2.21 gprocs

3.2.21.1 Inputs

The Emulator call gprocs shall take the following parameters.

ba: buffer address for process status information (in register R0)

3.2.21.2 Processing

The gprocs call shall copy process control information to a buffer in the calling user process domain. The process control data shall be at a lower or equal security level than the caller. The number of observed processes shall be

UNIX Emulator B5 Specification

returned. Only information about processes in the process family of the caller shall be returned. In the PDP-11/70 implementation the returned information shall adhere to the format of the existing UNIX process table structure to the extent possible. Security relevant information shall not be returned, and KSOS relevant information may be returned instead.

3.2.21.3 Outputs

The *Emulator call gprocs shall return the following values after execution.

np: number of process table entries returned to the user process
ec: error conditions

The Emulator call gprocs shall detect and report the following exception (error) conditions.

EX: bad buffer address

3.2.22 gtty

3.2.22.1 Inputs

The Emulator call gtty shall take the following parameters.

fd: open file descriptor
sb: pointer to data vector in user domain for status information

3.2.22.2 Processing

The Emulator call gtty shall return the current status of a device currently open and accessible for reading by the calling process.

3.2.22.3 Outputs

The Emulator call gtty shall return the following values after execution.

ec: error conditions
si: device status information

The Emulator call gtty shall detect and report the following exception (error) conditions.

EX: invalid file descriptor
EX: device is not a character mode device

3.2.23 indir

3.2.23.1 Inputs

The Emulator call indir shall take the following parameters.

sc: pointer to location containing a system call

3.2.23.2 Processing

The indir Emulator call shall execute the system call specified by the pointer argument. If indir is executed indirectly, no effect shall occur. Attempts to execute non-system calls shall result in a process fault.

3.2.23.3 Outputs

The Emulator call indir shall return the following values after execution.

ec: error conditions

The Emulator call shall detect and report the following exception (error) conditions.

EX: attempt to execute non-system call

EX: invalid address

3.2.24 kill

3.2.24.1 Inputs

The Emulator call kill shall take the following parameters.

pi: process ID of destination process

sn: number of signal to be sent

3.2.24.2 Processing

The Emulator call kill shall send the specified signal to the destination process. It shall not be possible for a process to kill itself. In the PDP-11/70 implementation of KSOS, the process number shall be passed in register R0. An error shall be detected and returned if the destination process has denied IPC writes via setting the discretionary access information for the process.

3.2.24.3 Outputs

The Emulator call kill shall return the following values after execution.

ec: error conditions

UNIX Emulator B5 Specification

The Emulator call kill shall detect and report the following exception (error) conditions.

- EX: destination process is the calling process
- EX: destination process has denied IPC writes to it (returned as if the process did not exist)
- EX: destination process does not exist
- EX: signal number is less than zero
- EX: signal number is greater than the Emulator maximum

3.2.25 link

3.2.25.1 Inputs

The Emulator call link shall take the following parameters.

- dn: pathname of file to which the link shall be made
- ln: pathname of the link

3.2.25.2 Processing

The link Emulator call shall create a link to a file with the specified pathname. The pathname interpretation and verification algorithm shall be applied to both specified pathnames. The full semantics of UNIX links shall be retained in KSOS.

3.2.25.3 Outputs

The Emulator call link shall return the following values after execution.

- ec: error conditions

The Emulator call link shall detect and report the following exception (error) conditions.

- EX: file system required for name interpretation is not mounted
- EX: file required for name interpretation is not a directory
- EX: directory on path is at a higher security level (returned as if the directory did not exist)
- EX: directory on path is at a lower integrity level (returned as if the directory did not exist)
- EX: directory on path has search permission denied
- EX: link name already exists
- EX: security level of directory to be written is higher than caller (returned as if there were no write permission on the directory)
- EX: integrity level of directory to be written is lower than caller (returned as if there were no write permission on the directory)
- EX: write permission to directory has been denied
- EX: attempt to link across file systems

UNIX Emulator B5 Specification

3.2.26 mknod

3.2.26.1 Inputs

The Emulator call mknod shall take the following parameters.

pn: pointer to new file pathname
da: discretionary access mode of new file
ad: "address" field (retained for UNIX compatibility)

3.2.26.2 Processing

The Emulator call mknod shall create a new file with the specified pathname. The discretionary access mode of the new file shall be set to the specified mode. (See the comments on chmod.) A directory shall (in all cases) be created. The creation of special (device) files shall be accomplished through system generation and NKSR Software functions. Unlike standard UNIX, regular user processes shall be able to create directories. Regular user processes shall not be able to write directly to directories as regulated by the Kernel subtype mechanism.

3.2.26.3 Outputs

The Emulator call mknod shall return the following values after execution.

ec: error conditions

The Emulator call mknod shall detect and report the following exception (error) conditions.

EX: insufficient integrity level to create a special file
EX: file in the searchable path is not a directory
EX: file system required for pathname interpretation is not mounted
EX: file exists with the same pathname

3.2.27 mount

The mount UNIX call has been replaced by the mount mechanism of the Non-Kernel Security-Related Software which has the same functionality.

3.2.28 nice

3.2.28.1 Inputs

The Emulator call nice shall take the following parameters.

np: new process priority (passed in R0)

3.2.28.2 Processing

The Emulator call nice shall advise the scheduler of a new runtime priority. This new priority shall only be an advisory. That is, it shall be the responsibility of the scheduler to guarantee equitable service. The legal values which the new priority may assume shall be greater than zero and less than some maximum, currently twenty (20.)

3.2.28.3 Outputs

The Emulator call nice shall return the following values after execution.

ec: error conditions

The Emulator call nice shall detect and report the following exception (error) conditions.

EX: new priority is less than zero

EX: new priority greater than system maximum

3.2.29 open

3.2.29.1 Inputs

The Emulator call open shall take the following parameters.

pn: pointer to pathname string of file to be opened

md: requested access mode { read, write, read and write }

3.2.29.2 Processing

The Emulator call open shall open the designated file for reading and/or writing. A file descriptor shall be returned for use in subsequent references to the file.

3.2.29.3 Outputs

The Emulator call open shall return the following values after execution.

fd: file descriptor (returned in register R0)

ec: error conditions

The Emulator call open shall detect and report the following exception (error) conditions.

EX: file required for name interpretation is not a directory

EX: file system required for name interpretation is not mounted

EX: directory on path is at higher security level (returned as if the directory did not exist)

UNIX Emulator B5 Specification

EX: directory on path is at lower integrity level (returned as if the directory did not exist)
EX: search permission on a pathname directory is denied
EX: file is at a higher security level and reading is requested (returned as if the file has read access denied)
EX: file is at a lower security level and writing is requested (returned as if the file had write access denied)
EX: file is at a lower integrity level and reading is requested (returned as if the file had read access denied)
EX: file is at a higher integrity level and writing is requested (returned as if the file has write access denied)
EX: discretionary permission is denied for writing, write requested
EX: discretionary permission is denied for reading, reading requested
EX: too many files are open

3.2.30 pipe

3.2.30.1 Inputs

The Emulator call pipe shall take the following parameters.

None

3.2.30.2 Processing

The Emulator call pipe shall create an input/output pipe. Two open file descriptors shall be returned, one for input and a second for output. In the PDP-11/70 implementation, the read file descriptor shall be returned in register R0. The write file descriptor shall be returned in register R1. Two available open file slots shall be required to successfully perform the pipe call.

3.2.30.3 Outputs

The Emulator call pipe shall return the following values after execution.

rf: read file descriptor
wf: write file descriptor

The Emulator call pipe shall detect and report the following exception (error) conditions.

EX: too many files open

3.2.31 port

3.2.31.1 Inputs

The Emulator call port shall take the following parameters.

UNIX Emulator B5 Specification

pn: port pathname string
mo: access mode for port
pf: pointer to a two integer block for returned file descriptors

3.2.31.2 Processing

(Not present in Version 6.0 UNIX.) Ports are a generalization of pipes. The port Emulator call shall create and open a new port. The specified pathname shall be either a complete pathname, a pathname relative to the process working directory, or the null string. Ports with the null string as their name shall be inaccessible by other processes. The port shall be at the same security and integrity level of the calling process. The mode parameter shall determine the format of the message headers attached at port data during write operations. The mode parameter shall also specify the write access to the port accorded to the creator, the creator's group and all other users. ([Zucker 77] and [Sunshine 77] discuss the supporting concepts and mechanization of ports in RAND UNIX.)

3.2.31.3 Outputs

The Emulator call port shall return the following values after execution.

fr: file descriptor for reading (also returned in register R0)
fw: file descriptor for writing

The Emulator call port shall detect and report the following exception (error) conditions.

EX: improper mode
EX: directory needed to interpret pathname is unmounted
EX: file on name interpretation path is not a directory
EX: search permission denied for directory on interpretation path
EX: attempted violation of mandatory security during name interpretation
 (returned as if directory had search permission denied)
EX: invalid pointer address for file descriptors
EX: port already exists

3.2.32 profil

3.2.32.1 Inputs

The Emulator call profil shall take the following parameters.

bp: pointer to histogram buffer in user domain
bs: buffer size (in bytes)
os: offset to be subtracted from sampled program counter
sc: scale to expand/contract observation

3.2.32.2 Processing

The Emulator call `profil` shall cause an execution profile to be compiled into a data vector in the user process domain. The user program counter will be sampled each 1/10 of a second thereafter while the process is physically executing. Offset and scale shall be used to adjust the program counter value to an index within the bounds of the histogram vector. If the resulting value is within the histogram vector, the observation is tallied. See [Thompson 75] for data types and formats.

In the PDP-11/70 implementation of KSOS, profiling shall be disabled by specifying a scale of zero or one. Profiling shall be rendered ineffective if the buffer size is zero. The execution of the call `exec` will disable profiling. Profiling shall remain enabled in both parent and child after a fork call.

3.2.32.3 Outputs

The Emulator call `profil` shall return the following values after execution.

ht: histogram tallies in the data vector
ec: error conditions

The Emulator call `profil` shall detect and report the following exception (error) conditions.

EX: improper buffer address in user process (not checked in 6.0 UNIX.)
EX: buffer size too large (not checked in 6.0 UNIX.)

3.2.33 ptrace

3.2.33.1 Inputs

The Emulator call `ptrace` shall take the following parameters.

rq: requested trace action
pi: process ID of traced (destination) process
ad: address parameter to action
da: data parameter to requested action

3.2.33.2 Processing

The Emulator call `ptrace` shall permit a parent process to control and monitor the execution of an immediate child process. Tracing of set-uid software shall not be permitted. The behavior of the `ptrace` call shall depend upon the interpretation of the requested action parameter. In the PDP-11/70 implementation of KSOS, the following action values and actions shall be supported.

- 0. executed by child. indicates that child expects to be traced.
- 1,2. get instruction space or data space word, respectively, from the child process.

UNIX Emulator B5 Specification

3. returned the addressed word from the per-process data table.
- 4,5. set word in the child instruction or data space, respectively.
6. write the per-process data word as addressed.
7. force a signal to the child process.
8. terminate the traced process.

With the exception of request zero (0), the traced process shall be stopped prior to execution of the tracing action. The wait call shall be used to determine when a process has stopped. Processes that have stopped shall return to the waiting parent the termination status value of 0177 (octal.)

Because ptrace may be implemented using the low bandwidth IPC mechanism, the usual security, integrity, and discretionary controls enforced by the Kernel may limit the extent of process tracing.

3.2.33.3 Outputs

The Emulator call ptrace shall return the following values after execution.

dv: data values
ec: error conditions

The Emulator call ptrace shall detect and report the following exception (error) conditions.

- EX: request value is out of bounds (less than 0 or more than 8)
- EX: requested action zero and parent does not exist
- EX: destination process does not exist
- EX: destination process has denied IPC writes to it (returned as if process did not exist)
- EX: destination process is not an immediate descendant of caller
- EX: requested data read or write and address was odd
- EX: requested data read or write and addressing error resulted

3.2.34 read

3.2.34.1 Inputs

The Emulator call read shall take the following parameters.

fd: file descriptor (passed in register R0)
ba: pointer to buffer for incoming data
nb: number of bytes to be read (if possible)

3.2.34.2 Processing

The read call shall transfer data from the specified file, device, or pipe to a buffer in the process user domain. The number of bytes specified in the call shall be transferred if possible. The actual number of bytes read shall be returned. In the PDP-11/70 implementation, the actual byte count shall be returned in register R0. The value zero shall be returned when end of file has

UNIX Emulator B5 Specification

been reached.

3.2.34.3 Outputs

The Emulator call read shall return the following values after execution.

db: data bytes to the user domain buffer
ac: actual count of data bytes transfered

The Emulator call read shall detect and report the following exception (error) conditions.

EX: invalid file descriptor
EX: file has not been opened for reading
EX: attempt to read from pipe without a writer at the other end
EX: physical device error
EX: illegal buffer address

3.2.35 seek

3.2.35.1 Inputs

The Emulator call seek shall take the following parameters.

fd: file descriptor (passed in register R0)
os: logical offset within file
pr: requested pointer adjustment action

3.2.35.2 Processing

The Emulator call seek shall adjust the logical file pointer for the designated open file. The pointer shall be adjusted according to the requested action as follows.

0: the pointer shall be set exactly to the parameter offset.
1: the pointer shall be set to the current seek pointer value plus the offset.
2: the pointer is set to the length of the file plus the offset.
3,4,5: the pointer shall be set as in 0,1,2 above but offset shall be multiplied by 512 before updating the pointer.
Absolute offset values shall be unsigned (case 0 and 3 above.) The offset shall otherwise be signed.

3.2.35.3 Outputs

The Emulator call seek shall return the following values after execution.

ec: error conditions

UNIX Emulator B5 Specification

The Emulator call seek shall detect and report the following exception (error) conditions.

EX: invalid file descriptor

EX: file descriptor refers to a pipe

EX: requested action is out of bounds (pr < 0 or pr > 5)

3.2.36 setgid

3.2.36.1 Inputs

The Emulator call setgid shall take the following parameters.

gi: new group id

3.2.36.2 Processing

The Emulator call setgid shall set the effective group id of the calling process. Processes which are appropriately privileged shall also set the real group id (i.e., the owner) of the process. Unprivileged processes shall be permitted to set the effective group id to the value of the real group id only.

3.2.36.3 Outputs

The Emulator call setgid shall return the following values after execution.

ec: error conditions

The Emulator call setgid shall detect and report the following exception (error) conditions.

EX: attempt to set effective group id to value other than real group id

3.2.37 setuid

3.2.37.1 Inputs

The Emulator call setuid shall take the following parameters.

nu: new effective user id

3.2.37.2 Processing

The Emulator call setuid shall set the effective user id of the calling process. Processes which are appropriately privileged shall also set the real user id (i.e., the owner) of the process. Unprivileged processes shall be permitted to set the effective user id to the value of the real user id only.

UNIX Emulator B5 Specification

3.2.37.3 Outputs

The Emulator call `setuid` shall return the following values after execution.

ec: error conditions

The Emulator call `setuid` shall detect and report the following exception (error) conditions.

EX: attempt to set effective user id to value other than real user id

3.2.38 sfork

3.2.38.1 Inputs

The Emulator call `sfork` shall take the following parameters.

None

3.2.38.2 Processing

The `sfork` call shall spawn a new process. The process creation function of this call shall be identical to the Emulator call `fork`. (See "fork" above.) The children generated by this call shall be immune from terminal generated interrupt (DEL character) or quit (FS character) signals. The children shall still receive direct signals from the kill Emulator call. (See "kill" above.)

3.2.38.3 Outputs

The Emulator call `sfork` shall return the following values after execution.

Identical to `fork`

The Emulator call `sfork` shall detect and report the following exception (error) conditions.

Identical to `fork`

3.2.39 signal

3.2.39.1 Inputs

The Emulator call `signal` shall take the following parameters.

sn: signal number

sf: signal function

UNIX Emulator B5 Specification

3.2.39.2 Processing

The Emulator call signal shall permit the user to specify the recovery action to be taken following the reception of a signal from another process. The PDP-11/70 implementation of KSOS shall provide for a minimum of thirteen (13) signals. The signal conditions shall be:

1. hangup
2. interrupt
3. quit (core image)
4. illegal instruction (not reset when caught - core image)
5. trace trap (not reset when caught - core image)
6. IOT instruction (core image)
7. EMT instruction (Kernel call - effect differs from standard UNIX)
8. floating point exception (core image)
9. kill (cannot be caught or ignored)
10. bus error (core image)
11. segmentation violation (core image)
12. bad argument to system call (core image)
13. write on a pipe with no one to read it

Note the difference from standard UNIX in signal number seven (7.) The user part of a process shall be able to directly call the Security Kernel. Because the EMT instruction shall be used to invoke the Kernel, the semantics of signal number seven must change. This signal slot shall be retained for historical reasons. Core dump images shall be generated for the signals so marked above.

The PDP-11/70 implementation of KSOS shall comply with the standard UNIX signal semantics relevant to the significance of signal function values and signal routine return procedures. If the signal function value is zero, the process shall terminate when the specified signal is received. This shall be the default action. If the signal function value is odd, the signal shall be ignored except where noted in the list above. An even valued signal function shall specify the address in the user domain of the signal handling routine. This routine shall be entered following the reception of the specified signal. A PDP-11/70 RTI or RTT instruction shall return from the signal interrupt. Signals which have been caught shall cause the selected action to be reset to zero. The number of signals to be set shall be a parameter to the Emulator generation procedure. The number shall not be less than thirteen.

3.2.39.3 Outputs

The Emulator call signal shall return the following values after execution.

ec: error conditions
ov: old signal value

The Emulator call signal shall detect and report the following exception (error) conditions.

EX: specified signal number is less than zero
EX: specified signal number is greater than the maximum number of signals permitted by the Emulator

UNIX Emulator B5 Specification

EX: specified signal is the kill signal (number 9)

3.2.40 sleep

3.2.40.1 Inputs

The Emulator call sleep shall take the following parameters.

st: sleep time in seconds

3.2.40.2 Processing

The Emulator call sleep shall suspend execution for the time period specified by the parameter.

3.2.40.3 Outputs

The Emulator call sleep shall return the following values after execution.

None

The Emulator call sleep shall detect and report the following exception (error) conditions.

None

3.2.41 stat

3.2.41.1 Inputs

The Emulator call stat shall take the following parameters.

pn: pointer to pathname string

ba: buffer address for file status information

3.2.41.2 Processing

The Emulator call stat shall return the file status information for the specified file. In the PDP-11/70 implementation of KSOS, the format of the information to be returned shall be consistent with the existing call except for such information which is meaningless under KSOS. These fields may be used for KSOS specific information. Certain values (e.g. link count, time of last access) may not be returned since they may be used as clandestine information channels. The access checks of the corresponding open call shall be satisfied in order for this call to succeed. If the file exists, but cannot be accessed due to mandatory security, integrity, or discretionary access reasons, a distinctive null status block shall be returned, however, the error condition shall not be raised.

UNIX Emulator B5 Specification

3.2.41.3 Outputs

The Emulator call stat shall return the following values after execution.

fs: file status information
ec: error conditions

The Emulator call stat shall detect and report the following exception (error) conditions.

EX: file system required for pathname interpretation is not mounted
EX: file required for pathname interpretation is not a directory
EX: directory on path is unsearchable
EX: directory on path is at higher security level and is unreadable
(returned as if directory was unsearchable)
EX: directory on path is at lower integrity level and is unreadable
(returned as if directory was unsearchable)
EX: file does not exist

3.2.42 stime

The stime UNIX call has been replaced by equivalent functionality in the Non-Kernel Security-Related Software.

3.2.43 stty

3.2.43.1 Inputs

The Emulator call stty shall take the following parameters.

fd: file descriptor
pp: pointer to parameter block
pb: device control parameter block

3.2.43.2 Processing

The Emulator call stty shall set the device dependent control parameters for the device referenced by the file descriptor. The format of the parameter block shall be the same as the existing call. Unlike Version 6.0 UNIX, the user must own the device for this call to be effective.

3.2.43.3 Outputs

The Emulator call stty shall return the following values after execution.

ec: error conditions

The Emulator call stty shall detect and report the following exception (error) conditions.

UNIX Emulator B5 Specification

EX: invalid file descriptor
EX: device is not a character mode device
EX: not owner of the device
EX: parameters invalid (stty can be used for other devices which may include additional checking on the parameter block)

3.2.44 sync

3.2.44.1 Inputs

The Emulator call sync shall take the following parameters.

None

3.2.44.2 Processing

The Emulator call sync shall cause all file system information for the process family to be updated.

3.2.44.3 Outputs

The Emulator call sync shall return the following values after execution.

ec: error conditions

The Emulator call sync shall detect and report the following exception (error) conditions.

None.

3.2.45 time

3.2.45.1 Inputs

The Emulator call time shall take the following parameters.

None

3.2.45.2 Processing

The Emulator call time shall return the system time in seconds expired since January 1, 1970. The system time shall be expressed in Greenwich Mean Time (GMT).

3.2.45.3 Outputs

The Emulator call time shall return the following values after execution.

UNIX Emulator B5 Specification

ti: time in GMT since January 1, 1970 (returned in R0 and R1)

The Emulator call time shall detect and report the following exception (error) conditions.

None

3.2.46 times

3.2.46.1 Inputs

The Emulator call times shall take the following parameters.

tb: pointer to data vector in user domain for execution time data

3.2.46.2 Processing

The times Emulator call shall return the time-accounting information for the currently executing process and its terminated children. The times shall be expressed in 1/10 seconds. The relevant information returned shall be the time spent by the process in user mode and Emulator mode execution, and the time spent by the terminated children in user mode and Emulator mode execution. The time elapsed in Kernel mode execution shall not be returned since this data opens a clandestine information channel through the Security Kernel.

3.2.46.3 Outputs

The Emulator call times shall return the following values after execution.

pu: process user mode execution time
ps: process Emulator mode execution time
cu: terminated child process user mode execution time
cs: terminated child process Emulator mode execution time
ec: error conditions

The Emulator call times shall detect and report the following exception (error) conditions.

EX: improper buffer address (not checked in Version 6.0 UNIX.)

3.2.47 umount


The UNIX u(n)mount call shall be replaced by equivalent functionality in the Non-Kernel Security-Related Software.

AD-A111 576


FORD AEROSPACE AND COMMUNICATIONS CORP PALO ALTO CA W--ETC F/8 9/2
K505 COMPUTER PROGRAM DEVELOPMENT SPECIFICATIONS (TYPE B-S). (K--ETC(U)
DEC 80
WDL-TR7932-REV-3
MDA903-77-C-0333
NL

UNCLASSIFIED

3 of 4
210106



11157



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

UNIX Emulator B5 Specification

3.2.48 unlink

3.2.48.1 Inputs

The Emulator call unlink shall take the following parameters.

pn: pointer to file pathname string

3.2.48.2 Processing

The unlink Emulator call shall remove the directory entry for a file. If this was the last reference (link) to the file, the file contents shall be deleted. If the file is currently open, the actual removal shall be delayed until the file has been closed. The directory entry shall be removed immediately.

3.2.48.3 Outputs

The Emulator call unlink shall return the following values after execution.

ec: error conditions

The Emulator call unlink shall detect and report the following exception (error) conditions.

- EX: file does not exist
- EX: file system needed to interpret pathname is not mounted
- EX: file on interpretation path is not a directory
- EX: directory on name interpretation path has search access denied
- EX: directory on interpretation path is at higher security level (returned as if directory had search access denied)
- EX: directory on interpretation path is at lower integrity level (returned as if directory had search access denied)
- EX: file to be removed is at lower security level (returned as if file had write access denied)
- EX: file to be removed is at higher integrity level (returned as if file had write access denied)
- EX: file to be removed has write permission denied
- EX: parent directory of file has write access denied
- EX: file to be removed is a directory and is not empty

3.2.49 wait

3.2.49.1 Inputs

The Emulator call wait shall take the following parameters.

None

UNIX Emulator B5 Specification

3.2.49.2 Processing

The Emulator call wait shall cause the executing process to suspend execution until the termination of one of its children. If no children exist, the return shall be immediate with the appropriate error code.

In the PDP-11/70 implementation of KSOS, the normal return shall present the process ID of the terminated child to the parent in R0. The high order byte of R1 shall contain upon return, the exit status of the process (see exit above.) The low order byte of R1 shall contain the termination status of the child process. The following values shall be used to represent specific termination circumstances.

0: normal termination.
0177: stopped, but not terminated. (see ptrace below.)
0200: terminated, core image produced.

3.2.49.3 Outputs

The Emulator call wait shall return the following values after execution.

pi: process ID of terminated child (returned in R0)
rs: child return status from its exit (R1 high order byte)
ts: child termination status (R1 low order byte)

The Emulator call wait shall detect and report the following exception (error) conditions.

EX: no children belonging to calling process

3.2.50 write

3.2.50.1 Inputs

The Emulator call write shall take the following parameters.

fd: file descriptor (passed in register R0)
ba: buffer address of data vector to be written
nb: number of bytes to be written

3.2.50.2 Processing

The write call shall transfer data from the process user domain to the specified file, device or pipe. The number of bytes actually written shall be returned. This count shall be returned in register R0, in the case of the PDP-11/70 KSOS implementation. If an error has resulted, the actual byte count shall differ from the requested transfer count.

UNIX Emulator B5 Specification

3.2.50.3 Outputs

The Emulator call write shall return the following values after execution.

bc: number of bytes actually written

ec: error conditions

The Emulator call write shall detect and report the following exception (error) conditions.

EX: invalid file descriptor

EX: file has not been opened for writing

EX: attempt to write a pipe without a reader at the other end

EX: physical input/output error

EX: illegal buffer address

EX: invalid requested transfer count

3.2.51 Directory Manager

The KSOS Emulator shall use the file subtype capability of the Security Kernel to implement UNIX compatible directories. For a detailed discussion of Kernel level file subtypes, the reader is referred to the KSOS Security Kernel Computer Program Development Specification (Type B5) [Kernel 78]. For this discussion, it is sufficient to state that the Kernel file subtype mechanism shall permit the overall set of file resources to be divided into subtype partitions. In order to access a file within a particular subtype partition, the process shall meet the discretionary access conditions predefined for files of that subtype, in addition to the usual mandatory security, integrity and discretionary access conditions. Access to files within a subtype shall be granted on the basis of the requested access mode (i.e., read, write, execute (or search)) and the access modes permitted to the requestor. The categories to which a user may belong shall be file subtype owner, user group, and the set of all users, conforming to the usual notion of UNIX discretionary access permissions.

The following paragraphs describe one possible implementation of the Directory Manager. The method presented below is a minimum privilege solution. It only requires the Directory Manager to temporarily execute in the discretionary access domain of the directory subtype owner and that the Directory Manager be privileged to use the K_link and K_unlink Kernel calls. These privileges are insufficient to allow the a Directory Manager at say, the SECRET level to create a directory entry in an UNCLASSIFIED directory. This would require that the Directory Manager also be privileged to violate the security *-property. Should the Government require that the Directory Manager be able to manipulate directories that are at security levels different than its (the Directory Manager's) current level, the Directory Manager shall be privileged to violate the security *-property, and shall control the bandwidth of the resulting storage channel. This privileged version of the Directory Manager shall also create the appropriate audit information whenever it actually violates the security *-property, and send the audit information via an LPC message to the Audit Capture Process [NKSRS 78].

UNIX Emulator B5 Specification

To support search-only directories, the directory manager must be able to read execute-only files. Therefore the directory manager shall be privileged to violate discretionary access in order to perform path name interpret functions on search only directories.

3.2.51.1 Inputs

KSOS Emulator files shall belong to one of two file subtypes, ordinary files and directory files. All user processes shall have read, write, and execute permission to ordinary files as defined by the subtype discretionary permissions. (Of course, before a user process can successfully access an ordinary file, the mandatory security and owner-permitted discretionary access conditions shall be fulfilled.) Directory files, however, shall have the following file subtype permissions:

- a. file subtype owner: read, write, execute (search) enabled
- b. subtype owner group: read, write, execute (search) enabled
- c. all other users: read and execute (search) enabled

All KSOS users shall be able to read and search directory files. Only the owner of the subtype (the Directory Manager) shall be able to write or delete a directory file. Since directory files are still subject to mandatory security and discretionary access constraints, directory searches shall fail when the appropriate access conditions have not been satisfied.

When a user process intends to perform a directory operation that requires writing on a directory file, the process shall create a child process using the Kernel primitive `K_spawn` and then shall wait for the child process to return. The `K_spawn` call shall specify that the child process shall be the Directory Manager (i.e. that the child shall begin executing the image of the Directory Manager). Through the argument segment, the process shall specify the requested directory operation, the pathname of the directory, and other necessary parameters (i.e., file pathname, file SEID.)

3.2.51.2 Processing

When invoked, the Directory Manager process shall execute in the discretionary access domain of the directory subtype owner. The Directory Manager shall obtain explicit access capability to the file subtype, directory, by performing a Kernel level open on the subtype SEID. The Directory Manager shall change its access domain to the real group and user ID of the invoking process. The Directory Manager shall open the specified directory file at the Kernel level supplying the open object descriptor returned from the open on the directory subtype. The Directory Manager shall then perform any one of the following operations as requested by the invoking process.

- a. create a new directory
- b. make an entry in an existing directory

c. delete an entry from an existing directory

d. delete an existing directory

e. perform path name interpret functions.

Error conditions or return values shall be sent to the invoking process via IPC message. The Directory Manager process shall exit and cause the invoking process (waiting upon the death of the child) to resume execution.

3.2.51.3 Outputs

The Directory Manager shall detect and report the following exception conditions:

EX: directory file does not exist

EX: could not interpret directory pathname due to security or discretionary access violation

EX: attempt to delete non-empty directory

EX: file could not be found in directory

Exception conditions shall be returned to the requesting process via IPC message.

In addition, when the Directory Manager is forced to violate the security *-property, the Directory Manager shall create an audit record of the event, and send it via an IPC message to the Audit Capture Process.

3.2.52 Network Interface

The Emulator shall support the interface to a modern computer network. The protocol shall be the standard DOD Internet and Transmission Protocol [Postel 80a] [Postel 80b]. Table I shows the allocation of TCP functions to the three CPCI's which make up KSOS. In the remainder of this paragraph the TCP functions performed by the Emulator are discussed.

UNIX Emulator B5 Specification

Table I
TCP Functional Allocation

	Security Kernel	UNIX Emulator	TCP Daemon
IMP Support	x		
Rendezvous			
Flow Control		x	
Data Stream			
Integrity:			
Checksum			x
Resequencing		x	
Retransmission		x	
Acknowledge		x	
Duplicate Detect		x	
Routing			x
Connection Alloc.			x
Urgent Processing		x	
Formatting		x	
Security Level of Message			x

3.2.52.1 Inputs

The inputs to the network interface are the data stream(s) to and from the network, and control information. The control information shall normally be supplied as a side effect of the UNIX calls to the network (e.g., open, close, read, and write). If additional control information is required, the stty and gtty user calls may be utilized. The interpretation of the UNIX open call mode argument shall be modified when the "file" being opened is a network logical connection. If the mode is an even number above some defined minimum, it shall be interpreted as the address of a control block in the user address space that defines the parameters required for a general network connection to be established. The exact format of this control block is left to the discretion of the development contractor. Modes 0, 1, and 2 shall utilize a default set of parameters. Odd mode values shall be considered to be in error.

3.2.52.2 Processing

3.2.52.2.1 Initiation of Connections

The Emulator shall initiate a connection by passing a shared segment name to the Network Daemon (part of the Non-Kernel Security-Related Software CPCI) using the K_post Security Kernel call. Note that if the security levels of the user and the network are such that the requested communication cannot be established, the K_post call will fail. This shared segment shall be used for all subsequent communication between the Emulator and the Network Daemon. The initial contents of the shared segment shall be the parameters needed by the Network Daemon to allocate the connection. The Network Daemon and the Emulator shall use a protocol between themselves that assures that the shared segment will not become corrupted.

3.2.52.2.2 Flow Control

The Emulator shall provide flow control on each of the connections from the user process. In TCP this consists of selective advancing the the windows. The Emulator shall be designed to allow it to be tuned to improve overall performance. The tuning may be a dynamic adjustment of the window updating or may be a static (parameter) adjustment.

3.2.52.2.3 Data Stream Integrity

The Emulator shall provide most of the functionality in the area of data stream integrity. This shall include acknowledgement of received octets (bytes), retransmission of unacknowledged transmissions, detection of duplicate octets in the received stream, and resequencing the streams. The Network Daemon shall perform the checksum function on incoming and outbound data.

3.2.52.2.4 Urgent Processing

(The subject of urgent processing is currently under intense debate in the TCP community. The design requirements presented here should be reviewed in light of the latest conclusions of the community.)

The Emulator should notify the user that "urgent" data has not yet been received by him (the user). The user should upon receipt of such notification read the data stream up to the urgent pointer. The development contractor may suggest the manner in which the user is notified of the existence of unprocessed urgent data, and the way in which the start of the urgent data is communicated to the user. The preferred mechanism shall be via the signal mechanism.

3.2.52.2.5 Formatting

The Emulator shall break up the outbound data stream into letters and segments. The Emulator may allow the user to specify that the end-of-letter (EOL) bit is to be set at some point in the stream. In general, the formatting of the stream into segments and letters should be transparent to the user program.

UNIX Emulator B5 Specification

3.2.52.3 Outputs

The output of the network interface function of the Emulator shall be the data streams to and from the user process and the network.

3.2.53 Special Requirements

The programming standards for the Emulator are discussed in the System Specification (Type A). Because the Emulator is not trusted by the Security Kernel, the emphasis on provability of the code is lessened.

If the Directory Manager is privileged (so it can create files in directories whose level is different than that of the process), it shall be subject to the same design and construction requirements as other privileged software.

3.2.53.1 Human Performance

This paragraph is not applicable to this specification.

3.2.53.2 Government Furnished Property List

The compiler for the language used to implement the Emulator may be furnished by the Government. The Government shall approve the implementation language to be used.

3.3 Adaptation

In general the Emulator should not require adaptation to the environment of a particular site. Because the Emulator is not trusted, it may be modified without affecting the security of KSOS. However, incorrect modification of the Emulator may adversely affect the UNIX interface. The parameters which control the sizes of buffer caches, and similar functions shall be adjustable by each site to tune the system.

4. QUALITY ASSURANCE PROVISIONS

4.1 Introduction

The underlying philosophy of KSOS quality assurance is to provide a convincing demonstration of the security and completeness of the system. Testing and quality assurance are a complement to the formal verification aspects of KSOS. All testing on KSOS shall be conducted at the development contractor's facility.

KSOS shall be subject to the technical reviews and audit procedures of MIL-STD-1521A, Appendices A-F. The following technical reviews and audits will be conducted:

- a. System Requirements Review (Appendix A)
- b. System Design Review (Appendix B)
- c. Preliminary Design Review (Appendix C)
- d. Critical Design Review (Appendix D)
- e. Functional Configuration Audit (Appendix E)
- f. Physical Configuration Audit (Appendix F)

The Critical Design Review shall include a review of any mathematical proofs showing that the design meets the requirements of the Government approved DoD security model.

The Functional Configuration Audit shall be the vehicle for the formal review of the design versus its specifications. The Physical Configuration Audit shall be the vehicle for the formal review of the "as built" computer programs versus both their design and their supporting documentation. In addition, the Physical Configuration Audit shall verify that the implementation requirements of the Type A System Specifications have been met. Both audits shall include review of any relevant mathematical proofs of correctness.

4.1.1 Category I Testing

The development contractor shall establish test plans to demonstrate that all the requirements of Section 3 have been met. These tests shall be designed to exercise all the interface options for each function. To the extent feasible, the tests shall exercise all combinations of options. The tests shall include sequences intended to fail, such as attempts to violate the system's security rules or to use unacceptable combinations of options. Government approval of all test plans is required.

The general mechanism for assuring that the requirements of Section 3 have been met shall be a comparison of the system "state" before and after the function has been invoked. To ensure repeatability all testing shall be automated if possible. Such automation may include (but is not limited to) shell scripts.

UNIX Emulator B5 Specification

test programs, and terminal emulation by another computer system.

4.1.2 Category II Testing

In addition to the tests described above, the entire KSOS system shall be subject to an overall system test. These Category II tests shall be carried out in accordance with the provisions of the development contract, but should include at least one test with at least ten (10) terminals and a network connection simultaneously active. The Category II tests should exercise the entire system by simulating a typical user load. The Category II testing should be automated as much as possible through the use of shell scripts, test load programs, and terminal emulation by another computer system.

Table II - Quality Assurance

Assurance techniques:

NA	Not Applicable
V	Verification
F	Formal Specification
T	Testing
D	Demonstration
A	Analysis
I	Inspection

Applicable test types:

1	Module level testing
2	Integration testing
3	Acceptance testing

Requirement	Assurance Techniques							Test types		
	NA	V	F	T	D	A	I	1	2	3
3.1.1.				x	x		x	x	x	x
3.1.1.1.				x	x		x	x	x	x
3.1.1.2.1.				x	x		x	x	x	x
3.1.1.2.2.				x	x		x	x	x	x
3.1.1.2.3.				x	x		x	x	x	x
3.1.1.2.4.				x	x		x	x	x	x
3.1.1.2.5.				x	x		x	x	x	x
3.2.1. - 3.2.50.				x	x	x	x	x	x	x
3.2.51.			*	x	x	x	x	x	x	x
3.2.52.			*	x	x	x	x	x	x	x
3.2.53.		x								
3.2.53.1.		x								
3.2.53.2.		x								
3.3.						x	x			

* Security related functions to be formally specified.

UNIX Emulator B5 Specification

5. PREPARATION FOR DELIVERY

This section is not applicable to this specification.

UNIX Emulator B5 Specification

6. NOTES

This section is not applicable to this specification.

UNIX Emulator B5 Specification

10. APPENDIX

SECURE MINICOMPUTER OPERATING SYSTEM (KSOS)
NON-KERNEL SECURITY RELATED SOFTWARE
COMPUTER PROGRAM DEVELOPMENT
SPECIFICATIONS (TYPE B5)

Department of Defense Kernelized Secure Operating System

Contract MDA 903-77-C-0333
CDRL 0002AH

Prepared for:

Defense Supply Service - Washington
Room 1D245, The Pentagon
Washington, DC 20310



**Ford Aerospace &
Communications Corporation**
Western Development
Laboratories Division

3939 Fabrian Way
Palo Alto, California 94303

NOTICE

The Department of Defense Kernelized Secure Operating System (KSOS) is being produced under contract to the U.S. Government. KSOS is intended to be compatible with the Western Electric Company's UNIX[®] Operating System (a proprietary product). KSOS is not part of the UNIX license software and use of KSOS is independent of any UNIX license agreement. Use of KSOS does not authorize use of UNIX in the absence of an appropriate licensing agreement.

This document is a major revision of an earlier draft dated July 1978. That and other previous editions are obsolete and should not be used.

UNIX and PWB/UNIX are trade/service marks of the Bell System.

DEC and PDP are registered trademarks of the Digital Equipment Corporation, Maynard MA.

TABLE OF CONTENTS

1.	SCOPE	1
1.1.	Identification	1
1.2.	Functional Summary	1
1.3.	Terminology	2
2.	APPLICABLE DOCUMENTS	5
2.1.	Government Documents	5
2.1.1.	Directives, Manuals and Standards	5
2.1.2.	Reports	6
2.2.	Non-Government Documents	6
2.3.	Other References Not Part of this Specification	6
3.	REQUIREMENTS	8
3.1.	Computer Program Definition	8
3.1.1.	Interface Requirements	8
3.1.1.1.	Interface Block Diagram	8
3.1.1.2.	Detailed Interface Definition	9
3.1.1.2.1.	Secure User Services	9
3.1.1.2.1.1.	Secure Initiator	9
3.1.1.2.1.2.	Secure Server	9
3.1.1.2.1.2.1.	Access Authentication Functions	10
3.1.1.2.1.2.1.1.	Processing Operator Requests	10
3.1.1.2.1.2.2.	Password Modification Function	11
3.1.1.2.1.2.3.	File Access Modification Function	11
3.1.1.2.1.2.4.	Level Preserving Copy Function	11
3.1.1.2.1.2.5.	Level Preserving Print Function	11
3.1.1.2.2.	System Operation Services	12
3.1.1.2.2.1.	Network Daemon	12
3.1.1.2.2.1.1.	IMP Interface	12
3.1.1.2.2.1.2.	Daemon Control Interface	12
3.1.1.2.2.1.3.	User Interface	13
3.1.1.2.2.2.	Line Printer Spooling	14
3.1.1.2.2.2.1.	User Interface	14
3.1.1.2.2.2.2.	Line Printer Daemon Control Interface	14
3.1.1.2.2.3.	Assign/Deassign	14
3.1.1.2.2.4.	Mount/Unmount	14
3.1.1.2.2.5.	System Startup/Shutdown	15
3.1.1.2.2.6.	Secure Mail	15
3.1.1.2.3.	System Maintenance Services	15
3.1.1.2.4.	System Administrative Services	15
3.1.1.2.4.1.	User Control Editor	15
3.1.1.2.4.2.	Immigration Officer Process	16
3.1.1.2.4.3.	Audit Capture Process	16
3.1.1.2.4.4.	Privilege Control Process	16
3.2.	Detailed Functional Requirements	17
3.2.1.	Secure User Services	17
3.2.1.1.	Secure Initiator	17

3.2.1.1.1.	Users and Privileges	17
3.2.1.1.2.	Inputs	17
3.2.1.1.3.	Processing	17
3.2.1.1.4.	Outputs	18
3.2.1.2.	Secure Server	18
3.2.1.2.1.	Users and Privileges	18
3.2.1.2.2.	Inputs	18
3.2.1.2.3.	Processing	18
3.2.1.2.4.	Outputs	18
3.2.1.3.	Access Authentication Functions	19
3.2.1.3.1.	Inputs	19
3.2.1.3.2.	Processing	19
3.2.1.3.2.1.	Login Processing	19
3.2.1.3.2.2.	Changing Group	20
3.2.1.3.2.3.	Change Access Level	20
3.2.1.3.2.4.	Logout	20
3.2.1.3.2.5.	Processing Operator Requests	20
3.2.1.3.3.	Outputs	20
3.2.1.4.	Password Modification Function	21
3.2.1.4.1.	Inputs	21
3.2.1.4.2.	Processing	21
3.2.1.4.3.	Outputs	21
3.2.1.5.	File Access Modification Function	21
3.2.1.5.1.	Inputs	22
3.2.1.5.2.	Processing	22
3.2.1.5.3.	Outputs	22
3.2.2.	System Operation Services	22
3.2.2.1.	Network Daemon	22
3.2.2.1.1.	Users and Privileges	23
3.2.2.1.2.	Inputs	23
3.2.2.1.3.	Processing	23
3.2.2.1.3.1.	Rendezvous	24
3.2.2.1.3.2.	Checksum	24
3.2.2.1.3.3.	Routing	24
3.2.2.1.3.4.	Connection Allocation	24
3.2.2.1.3.5.	Message Security Level	24
3.2.2.1.4.	Outputs	25
3.2.2.2.	Line Printer Spooling	25
3.2.2.2.1.	Users and Privileges	25
3.2.2.2.2.	Inputs	25
3.2.2.2.3.	Processing	25
3.2.2.2.4.	Outputs	27
3.2.2.3.	Assign/Deassign	27
3.2.2.3.1.	Users and Privileges	27
3.2.2.3.2.	Inputs	27
3.2.2.3.3.	Processing	27
3.2.2.3.4.	Outputs	28
3.2.2.4.	Mount/Unmount	28
3.2.2.4.1.	Users and Privileges	28
3.2.2.4.2.	Inputs	28
3.2.2.4.3.	Processing	28
3.2.2.4.4.	Outputs	29
3.2.2.5.	System Startup/Shutdown	29
3.2.2.5.1.	Users and Privileges	29

3.2.2.5.2.	Inputs	29
3.2.2.5.3.	Processing	29
3.2.2.5.4.	Outputs	30
3.2.2.6.	Secure Mail	30
3.2.2.6.1.	Users and Privileges	30
3.2.2.6.2.	Inputs	30
3.2.2.6.3.	Processing	30
3.2.2.6.4.	Outputs	30
3.2.3.	System Maintenance Services	30
3.2.3.1.	File-system Backup and Restoration	31
3.2.3.1.1.	Dump	31
3.2.3.1.1.1.	Users and Privileges	31
3.2.3.1.1.2.	Inputs	31
3.2.3.1.1.3.	Processing	31
3.2.3.1.1.4.	Outputs	31
3.2.3.1.2.	Restore	31
3.2.3.1.2.1.	Users and Privileges	31
3.2.3.1.2.2.	Inputs	32
3.2.3.1.2.3.	Processing	32
3.2.3.1.2.4.	Outputs	32
3.2.3.2.	File System Maintenance Services	32
3.2.3.2.1.	Make File System	32
3.2.3.2.1.1.	Users and Privileges	32
3.2.3.2.1.2.	Inputs	32
3.2.3.2.1.3.	Processing	32
3.2.3.2.1.4.	Outputs	32
3.2.3.2.2.	Directory Consistency Check	33
3.2.3.2.2.1.	Users and Privileges	33
3.2.3.2.2.2.	Inputs	33
3.2.3.2.2.3.	Processing	33
3.2.3.2.2.4.	Outputs	33
3.2.3.2.3.	Storage Consistency Check	33
3.2.3.2.3.1.	Users and Privileges	33
3.2.3.2.3.2.	Inputs	33
3.2.3.2.3.3.	Processing	33
3.2.3.2.3.4.	Outputs	33
3.2.3.2.4.	Kernel Name to Pathname Mapping	34
3.2.3.2.4.1.	Users and Privileges	34
3.2.3.2.4.2.	Inputs	34
3.2.3.2.4.3.	Processing	34
3.2.3.2.4.4.	Outputs	34
3.2.3.2.5.	Modify Control Entries for Filesystem	34
3.2.3.2.5.1.	Users and Privileges	34
3.2.3.2.5.2.	Inputs	34
3.2.3.2.5.3.	Processing	34
3.2.3.2.5.4.	Outputs	34
3.2.3.2.	KSOS System Generation	35
3.2.3.3.1.	Users and Privileges	35
3.2.3.3.2.	Inputs	35
3.2.3.3.3.	Processing	35
3.2.3.3.1.	Bootstrap	35
3.2.3.3.3.2.	Starter System	36
3.2.3.3.3.3.	Libraries	36
3.2.3.3.3.4.	Special Cases	36
3.2.3.3.3.5.	System Generation	36
3.2.3.3.3.6.	Integrity of the Inputs	37

3.2.3.3.3.7.	Rename System	37
3.2.3.3.3.8.	System Regeneration	37
3.2.3.3.4.	Outputs	37
3.2.4.	System Administration Services	37
3.2.4.1.	User Control Editor	37
3.2.4.1.1.	Users and Privileges	38
3.2.4.1.2.	Inputs	38
3.2.4.1.3.	Processing	38
3.2.4.1.4.	Outputs	38
3.2.4.2.	Audit Capture Process	38
3.2.4.2.1.	Users and Privileges	38
3.2.4.2.2.	Inputs	38
3.2.4.2.3.	Processing	39
3.2.4.2.3.	Outputs	39
3.2.4.3.	Privilege Control Process	39
3.2.4.3.1.	Users and Privileges	39
3.2.4.3.2.	Inputs	39
3.2.4.3.3.	Processing	39
3.2.4.3.4.	Outputs	39
3.2.4.4.	Immigration Officer Process	39
3.2.4.4.1.	Users and Privileges	40
3.2.4.4.2.	Inputs	40
3.2.4.4.3.	Processing	40
3.2.4.4.4.	Outputs	40
3.2.5.	Special Requirements	40
3.2.5.1.	Human Performance	40
3.2.5.2.	Government Furnished Property List	40
3.3.	Adaptation	40
4.	QUALITY ASSURANCE PROVISIONS	41
4.1.	Introduction	41
4.1.1.	Category I Testing	41
4.1.2.	Category II Testing	42
5.	PREPARATION FOR DELIVERY	45
6.	NOTES	46
10.	APPENDIX A	47

Kernelized Secure Operating System (KSOS): Non-
Kernel Security Related Software Computer
Program Development Specification (Type B5)

1. SCOPE

1.1 Identification

This specification establishes the performance, design, development and test requirements for the Non-Kernel Security Related Software Computer Program Configuration Item (CPCI) which is part of the Kernelized Secure Operating System (referred to as "KSOS"). KSOS provides a provably secure, resource-sharing operating system compatible with the standard user environment provided by UNIX, as described by [Thompson 75] and [Ritchie 74]. The Non-Kernel Security Related Software provides those functions which are needed by the system for continued operation and maintenance, and which are not part of the Security Kernel CPCI.

1.2 Functional Summary

The Non-Kernel Security Related Software is a collection of autonomous subsystems; some of which execute with extraordinary privilege (i.e. more privilege than is afforded to user programs) to provide essential services to the system. Such services include, but are not limited to, the following:

- ◆ Secure User Services; those services invoked by users which require a trusted path from the user to the service (e.g., the login process).
- ◆ System Operation Services; those functions essential to the operation of a KSOS system (e.g., the Process Bootstrapper).
- ◆ System Maintenance Services; those functions needed for continued operation and maintenance of a KSOS system, such as dump and restore of file systems.
- ◆ System Administrator Services; those functions needed to support the administrative operation of the system, such as the adding and deleting of users.

The remainder of this specification is organized as follows:

- ◆ Section 2 contains the references
- ◆ Section 3 contains the requirements for the CPCI
- ◆ Section 4 contains the quality assurance provisions for the CPCI
- ◆ Sections 5 and 6 are not applicable and are, therefore not present.

Non-Kernel Security Related Software B5 Specification

1.3 Terminology

The language used throughout this specification attempts to conform to the guidelines of Section 3.2.3 of MIL-STD-490. In particular, the word "shall" means that the specification expresses a provision that is binding. The words "should" and "may" mean that the specification expresses a provision which is non-mandatory. The word "will" is used to express a declaration of purpose on the part of the Government.

KSOS shall support the mandatory DoD security policy of DoD Directive 5200.1-R which is embodied in a Government approved mathematical model. [Verif 78] Proofs of the system's security properties shall be in terms of this model. KSOS shall provide a minimum of eight (8) hierarchical security classifications categories (or simply security categories) and a minimum of thirty-two (32) mutually independent security compartments. The security categories shall be such that

UNCLASSIFIED < CONFIDENTIAL < SECRET < TOP SECRET

Where "<" is defined in accordance with the requirements of DoD 5200.1-R. One security compartment shall be reserved for read protection of system data bases and programs. This compartment shall be called the "system" compartment.

KSOS shall provide for Kernel-enforced integrity. Integrity is defined as the formal mathematical dual of security [Biba 75]. At least four (4) hierarchical integrity classifications categories (or simply integrity categories) shall be provided in KSOS. The integrity categories include at least the following three classification categories:

USER < OPERATOR < ADMINISTRATOR

Although there is no immediate requirement for integrity compartments, the development contractor should include provisions to ease their later inclusion.

For the remainder of this specification, the term "security level" means the combination of a security category and a set of security compartments (which may be null). The term "integrity level" means the combination of an integrity category, and a (presently always null) set of integrity compartments. The term "level" (or access level) means the security and integrity level; i.e., the combination of a security category, a set of security compartments (which may be null), an integrity category, and a (presently always null) set of integrity compartments. The "access matrix" defines, for a particular user or group, the combination of the maximum security level permitted for that user or group, all security compartments to which that user or group has access, the maximum

Non-Kernel Security Related Software B5 Specification

integrity level for that user or group, and all integrity compartments (presently always null) to which that user or group has access.

In a secure system, it is necessary to have processes external to the kernel which perform security critical functions. These processes have come to be called "trusted processes". In KSOS, there is a refinement in this terminology. As Figure 1-1 shows there are subdivisions of the software inside the security perimeter. Some of the processes are "privileged", in other words, they may violate one or more of the security rules enforced by the Kernel. These processes must be designed and implemented with the same care and thoroughness as the Kernel itself. Section 3.1.1.2.4.4 presents the privileges which may be afforded processes.

A second subdivision has been termed "responsible" processes. These processes are not privileged to violate any of the Kernel's rules, but which are, nonetheless, important to the overall security of the system. Examples of responsible processes include the Secure Mail facility and the processes which manipulate the data bases describing the users' security and integrity permissions (the "access matrix"). These processes must also be carefully designed and implemented, but do not require formal specification or verification. Because they do perform security related functions, they cannot include untrusted components (e.g. the UNIX Emulator).

The third subdivision is the "encapsulated utilities" (EU in the figure). These are self-contained functions to which the user simply supplies parameters. The encapsulated utilities are not privileged to violate Kernel rules and do not affect the security of users. The chief function in the encapsulated utilities is system generation.




Figure 1-1

Non-Kernel Security Related Software B5 Specification

2. APPLICABLE DOCUMENTS

The following documents, of exact issue shown, form a part of this specification to the extent specified herein. In the event of a conflict between the referenced documents and the contents of this specification, this specification shall be considered a superseding requirement. In the text, references to these documents are in the form [Name date], e.g. [Biba 75].

2.1 Government Documents

2.1.1 Directives, Manuals and Standards

- a. DoD 5200.1-R Information Security Program Regulation
- b. DoD 5200.28 Security Requirements for Data Processing (ADP)
- c. DoD 5200.28-M ADP Security Manual
- d. MIL-STD-483 Configuration Management
- e. MIL-STD-490 Specification Practices
- f. MIL-STD-1521A Technical Reviews and Audits

2.1.2 Reports

- a. [A-Specs 78] "KSOS System Specification (Type A)", WDL-TR7808 Revision 1, Ford Aerospace and Communications Corporation, Palo Alto, CA (July 1978).
- b. [Bell and LaPadula 73] Bell, D.E. and LaPadula, L.J., "Secure Computer Systems", ESD-TR-73-278, Volume I-III, MITRE Corporation, Bedford, MA (November 1973 - June 1974).
- c. [Biba 75] Biba, K.J., "Integrity Considerations for Secure Computer Systems", MTR-3153, MITRE Corporation, Bedford, MA (June 1975).
- d. [Emulator 78] "KSOS Computer Program Development Specification (Type B5): UNIX Emulator", WDL-TR7932, Ford Aerospace and Communications Corporation, Palo Alto, CA (September 1978).
- e. [Kernel 78] "KSOS Computer Program Development Specification (Type B5): Security Kernel", WDL-TR7932, Ford Aerospace and Communications Corporation, Palo Alto, CA (September 1978).
- f. [Verif 78] "KSOS Verification Plan", WDL-TR7809, Ford Aerospace and Communications Corporation, Palo Alto, CA (March 1978).
- g. [Walter et al. 74] Walter, K.G. et al., "Primitive Models for Computer Security", ESD-TR-74-117, Case Western Reserve University, Cleveland, OH (January 1974).

Non-Kernel Security Related Software B5 Specification

2.2 Non-Government Documents

- a. [BBN 75] "Interface Message Processor, Specifications for the Interconnection of a Host and an IMP", Report 1822, Bolt Beranek and Newman Inc., Cambridge, MA, (December 1975).
- b. [Holmgren et al. 77] Holmgren, S.F., Healy, D.C., Jones, P.B. and Kasprzycki, E., "Illinois Inter-Process Communication Facility for UNIX", CAC Technical Memorandum 84, CCTC-WAD Document 7507, Center for Advanced Computation, University of Illinois at Urbana-Champaign, Urbana, IL (April 1977).
- c. [Lampson 73] Lampson, B., "A Note on the Confinement Problem", CACM, Volume 16, Number 10, pp 613 - 615 (October 73).
- d. [Lipner 75] Lipner, S.B., "A Comment on the Confinement Problem", Proc. Fifth Symposium on Operating Systems Principles, ACM SIGOPS Review, Volume 9, Number 5, pp 192 - 196 (19-21 November 1975).
- e. [Nemeth et al. 77] Nemeth, A.G., Sunshine, C., Zucker, S. and Tepper, "Progress Towards a DeFacto DoD UNIX Inter-Process Communication Standard", Working Paper, (May 1977). (Available from the authors.)
- f. [Parnas 72] Parnas, D.L., "A Technique for Software Module Specification with Examples", CACM, Volume 15, Number 5, pp 330 - 336 (May 1972).
- g. [Postel 78a] Postel, J.B., "Internetwork Protocol Specification", Version 4, Information Sciences Institute, University of Southern California, Marina del Rey, CA (February 1979).
- h. [Postel 78b] Postel, J.B., "Specification of Internetwork Transmission Control Protocol - TCP", Version 4 (draft), Information Sciences Institute, University of Southern California, Marina del Rey, CA (February 1979).
- i. [Roubine and Robinson 77] Roubine, O., L. Robinson, Special Reference Manual, 3rd ed., Technical Report CSG-45, SRI International, Menlo Park, CA (January 1977).
- j. [Sunshine 77] Sunshine, C., "Interprocess Communication Extensions for the UNIX Operating System: I Design Considerations", R-2064/1-AF, RAND Corporation, Santa Monica, CA (June 1977).
- k. [Zucker 77] Zucker, S., "Interprocess Communication Extensions for the UNIX Operating System: II Implementation", R-2064/2-AF, RAND Corporation, Santa Monica, CA (June 1977).

2.3 Other References Not Part of This Specification

- a. [Thompson 75] Thompson, K. and Ritchie, D.M., "UNIX Programmer's Manual", Sixth Edition, Western Electric Corporation, Greensboro, NC (May 1975).
- b. [Ritchie 74] Ritchie, D.M. and Thompson, K., "The UNIX Timesharing System", CACM, Volume 17, Number 5, pp 365 - 375 (May 1974).

Non-Kernel Security Related Software B5 Specification

- c. [UNIX 75] "UNIX Program Listings", Version 6.0, Western Electric Corporation, Greensboro, NC (May 1975).

Non-Kernel Security Related Software B5 Specification

3. REQUIREMENTS

The Non-Kernel Security Related Software shall provide those security relevant functions which are not allocated to the Security Kernel CPCI. Certain functions of the CPCI may be privileged to violate one or more of the security rules of the Security Kernel in providing overall system security. These privileged functions and the rules which they violate are identified in this section. The privileged functions must be implemented with the same precision and thoroughness as the Security Kernel. The System Specifications (Type A) [A-Specs 78] for KSOS present the requirements and constraints for the design and the implementation. Non-Kernel Security Related Software shall operate correctly without modification on all hardware configurations defined in the System Specification for KSOS.

Several of the functions described in this section require constraints on their user; e.g., to use some of these programs, the user must be at a certain integrity level or higher. These programs must be run in the "SYSTEM" compartment which is reserved for such programs and data bases. These steps assure that the executable images and data bases remain untampered with and that the programs cannot be invoked successfully by unauthorized users. For compactness, some of the functions are described as being available only to the System Operator, which means a user running at OPERATOR level or higher level and in the "system" security compartment.


3.1 Computer Program Definition

3.1.1 Interface Requirements

The Non-Kernel Security Related Software may utilize the Security Kernel interface and, in some cases, may use the UNIX interface provided by the UNIX Emulator. These two interfaces are specified in the companion Type B5 specifications for these CPCI's.

3.1.1.1 Interface Block Diagram

Figure 3-1 shows the relationship of the functional components of KSOS.



(NKSR: Non-Kernel Security Related Software)

Figure 1-2. Functional Components of KSOS

Non-Kernel Security Related Software B5 Specification

The call from domains of lesser privilege to those of greater privilege shall be mechanized by means of an inter-domain call. In the PDP-11/70 implementation the TRAP and IOT instructions shall be used by user programs to call the UNIX Emulator. The EMT instruction shall be used by user programs and the Emulator to call upon the Kernel.

It should be noted that no privileged functions may use the Emulator, since the Emulator is not trusted. Those functions in the Non-Kernel Security Related Software which are not privileged (e.g., the Secure Mail spooler) may use the Emulator.

3.1.1.2 Detailed Interface Definition

3.1.1.2.1 Secure User Services

The Secure User Services class of Non-Kernel Security Related Software consists of those functions which require a secure path to and from the user, such as the login process. In general, the functions which require such a secure path are those in which the user may be supplying raw passwords or the names of objects to be manipulated in a trusted way such as changing the security of an object.

3.1.1.2.1.1 Secure Initiator

The Secure Initiator shall be invoked by the Initial Process when the system is booted. The Secure Initiator shall spawn off the Audit Capture Process (ACP), initialize various objects, and execute a file of commands similar to the rc files on UNIX. The Initiator shall also spawn one Secure Server process (SSP) for each configured terminal in the system. A data base containing a list of the configured terminals paired with the appropriate Secure Servers id created by the Secure Initiator.

3.1.1.2.1.2 Secure Server

One Secure Server process shall be present for each configured terminal in a KSOS system. The Secure Server shall receive commands from the user at the terminal via the secure path, and shall spawn processes on behalf of the user to carry out these commands.

Each terminal in the system shall logically consist of a set of mutually exclusive logical data-paths. Only one path shall be active at any one time. All terminal input shall be directed to the currently active path, and an attempt to read or write on a path other than the active path shall cause the requesting process to become blocked. One path shall be denoted as the secure path, and shall be used exclusively by the Secure Server. The secure attention character shall cause the kernel to select the secure path as the active path. Selection of an alternative terminal path shall be performed by a SET_PATH operation of the K_device_function call. This shall be a privileged operation whose use is restricted to the Secure Server.

On receipt of a secure attention character, the Secure Server shall prompt for input from the terminal. The Secure Server shall interpret the input as a request for a Secure User Service. If the user is authorized to use the service,

Non-Kernel Security Related Software B5 Specification

the service shall be performed, either by the Server, or by spawning a process to perform the service on behalf of the user.

Some Secure User Services shall be capable of changing their security and integrity levels, and real owner (user ID and group ID) while running. This capability provides the mechanism necessary to perform in a controlled manner those user requests which either change the user's level or which require actions to be performed at another level. Because the Server has access to the user's access matrix it is able to verify the user's access level restrictions before performing any requested functions. Additionally, certain Secure processes can access restricted databases which system users are unable to access. These processes allow for the updating of these databases in a controlled and audited manner.

The Secure User Services under the direction of the Secure Server shall provide for authentication of users, control of user access levels, password modification, file access modification, and the spawning of level preserving copy and print functions.

3.1.1.2.1.3 Access Authentication Functions

A set of the Secure User Services known as the Access Authentication Functions shall service all user requests for login, changing group, changing user access level (access level includes security and integrity level), and logout.

3.1.1.2.1.3.1 Processing Operator Requests

The login and logout functions shall accept the following commands from the System Operator to control system access.

- a. Logout a Terminal - Requests to logout a terminal shall be handled in the same manner as if requested by the user. Whether or not the terminal was logged in, the terminal mode shall be set to the values listed for it in the Terminal Profile Data Base. This provides a mechanism to regain control of a terminal in an incorrect mode.
- b. Logout All Terminals - A check shall be made of all terminals configured in the Terminal Profile Data Base and a logout shall be performed on all terminals except the terminal designated as the operator's console.
- c. Lock a Terminal - This command shall prevent logins on the designated terminal. This function will not force a logout on a terminal that is logged in when the command is issued, although the System Operator may do so with another command if desired.
- d. Lock All Terminals - This command shall prevent login on all terminals except the operator's console. However, terminals already logged in will not be logged out.
- e. Unlock a Terminal - This command shall allow logins on the designated terminal.

Non-Kernel Security Related Software B5 Specification

- f. Unlock All Terminals - This command shall allow logins on all terminals.
- g. Set System Security Level - This command shall set the maximum security level permitted on the system and this maximum level shall be checked when a user makes a change level or login request. It shall not affect the level of any user logged in when the command is issued. If the system level is being lowered, the System Operator may wish to logout users above that level, or he (she) may simply notify them to change their levels or logout as soon as possible, forcing a logout only if they are unresponsive. It shall not be possible to raise the system level above the system maximum level defined in the System Database.

3.1.1.2.1.4 Password Modification Function

The Password Modification Function shall allow a user to change his login password. It shall also allow the user designated as the Group Administrator in the Group Access Authentication Database to change the group password for that group. The Password Modifier shall not allow a user to change the password of another user. If a user is unable to login because of a bad password he must request the System Administrator to change it to a temporary password, allowing him to login and enter a new password.

3.1.1.2.1.5 File Access Modification Function

The File Access Modification Function shall allow a user to modify access levels including security, integrity, and discretionary access levels on files which he owns. It shall not allow him to grant any special privileges to his files.

3.1.1.2.1.6 Level Preserving Copy Function

The Secure User Services shall provide facilities to copy files while maintaining their security and integrity levels without the requirement of the user being at the same level as the file. Since objects created by processes are always created at the access level at which the process is running, normal copies would often result in the unnecessary upgrading of files. All security checks shall remain in effect for this type of copy. The test to determine if a Level Preserving Copy will be successful is to determine if the user would be able to change his level to the level of the file and then attempt do a normal copy.

If the above test is successful the Secure Server shall spawn a standard copy at the level of the file to be copied and with the requesters owner and group. The copy may still fail if the copy process is unable to open the input file or create the output file because discretionary access or other constraints.

3.1.1.2.1.7 Level Preserving Print Function

The Level Preserving Print Function operates in the same manner as the Level Preserving Copy Function except that the output is spooled for printing on the line printer.

3.1.1.2.2 System Operation Services

The System Operation Services class of Non-Kernel Security Related Software shall consist of those functions which are essential to the continuing operation of the system. These functions shall have separate interfaces and, in general, are independent of each other.

3.1.1.2.2.1 Network Daemon

The Network Daemon shall multiplex and demultiplex the data streams to and from the computer network. The Network Daemon shall be privileged to allow it to handle a network that can support communications at varying security levels. The Network Daemon shall have three main interfaces:

- a. the Interface Message Processor (IMP) (or equivalent) interface
- b. the Daemon Control Interface
- c. the Network User Interface

3.1.1.2.2.1.1 IMP Interface

The Network Daemon shall communicate with the network via the IMP interface that is provided by the Security Kernel. This IMP interface shall appear to the Network Daemon as a device. For functions not common to all devices, the K_device_function Kernel call shall be employed. The main part of the IMP interface of the Network Daemon shall be the mechanization of the Host-to-Imp Protocol as defined by [BBN 75] (or other equivalent interfaces as specified by the Government in the development contract) and the mechanization of the Internet-work Datagram Layer [Postel 78a] of the Transmission Control Protocol (TCP).

3.1.1.2.2.1.2 Daemon Control Interface

The Daemon Control Interface shall be the mechanism by which the System Operator (or other authorized user) controls the state of the Network Daemon, and thereby the state of all network connections. The Daemon Control Interface shall be mechanized by means of inter-process communication from another process running at the OPERATOR or higher integrity level. The following Daemon Control commands shall be provided (the development contractor may extend or modify this list to suit the needs of a particular Network Daemon implementation):

- a. Reset Daemon - shall return the Network Daemon to its initial state. This function shall include destroying all logical connections, transmitting any required resetting messages to remote hosts, and similarly informing the local end of all network connections. It should be used only in cases where the Network Daemon has become hopelessly confused.
- b. Reset Host - shall return the state of a particular remote host to its initial state, destroying all open connections to that host.
- c. Abort Connection - shall destroy a particular connection. Both the local and remote ends of the connection shall be informed of the aborting of the connection via the normal Network Daemon status returns.

Non-Kernel Security Related Software B5 Specification

- d. Dump Status - shall cause the Network Daemon to dump its internal status information to a file. Because the status of the Daemon may involve connections of differing security levels, the classification of the status file shall be system high.

3.1.1.2.2.1.3 User Interface

The User Interface to the Network Daemon shall be via inter-process communication and shared segments. The bulk of the TCP functions shall be performed in the UNIX Emulator [Emulator 78]. Thus, a UNIX user-mode program would not communicate directly with the Network Daemon. Rather, the system calls from the user-mode program shall be interpreted by the UNIX Emulator and transformed into inter-process communication messages to and from the Network Daemon. The user shall initiate network communications using the UNIX "open" system call. To provide the additional information needed to open a network connection, the "mode" argument shall be subject to an extended interpretation. If the "mode" argument is even and greater than some defined constant and the "open" is for a network connection, then the argument shall be interpreted as the address of a control block. The exact format of the control block is left to the discretion of the development contractor.

The UNIX Emulator shall interface with the Network Daemon in the following manner. The opening of connections shall be accomplished by the transmission of IPCs to the daemon. The IPCs shall specify the identity of a sharable segment for subsequent data transfers between the emulator and the daemon. Completion of a connection open request by the daemon shall be notified to the emulator by an IPC message containing status information. Reading to and writing from the network shall be accomplished by passing the data via the shared segment, with control information for each transfer being passed in IPC messages. Multiple connections from a particular user may be multiplexed over the same shared segment. The UNIX Emulator and the Network Daemon shall obey a protocol that preserves the integrity of the shared segment (for example, the protocol could ensure that they both do not try to write the segment simultaneously, although other less restrictive protocol possibilities exist).

To facilitate the construction of network applications which do not include the UNIX Emulator, the TCP handling in the Emulator shall be designed to be as separate as possible from the remainder of the Emulator.

3.1.1.2.2.2 Line Printer Spooling

KSOS shall provide a secure mechanism for spooling output to the printer. The mechanism shall be secure in the sense that users cannot alter, view, or delete the spooled output of other users. In addition, the appropriate headers and footers shall be placed on the printed output. Finally, a record of the creation of classified printer output shall be generated. There are two interfaces to the Spooler:

- a. The User Interface - similar to the interface to the present UNIX spooler (/bin/opr and /lib/lpr).
- b. The Line Printer Daemon Control Interface - allows a System Operator (or other privileged user) to control the Daemon (de-spooler).

Non-Kernel Security Related Software B5 Specification

3.1.1.2.2.2.1 User Interface

The interface to the Line Printer Spooler shall allow the user to:

- a. Print multiple copies of a file (or files)
 - b. Allow the user to specify an alternate heading on the burst pages
- The User Interface shall be implemented as arguments to the Spooler program, using the standard UNIX argument passing conventions.

3.1.1.2.2.2.2 Line Printer Daemon Control Interface

The Line Printer Daemon Control Interface shall be via inter-process communication and shall allow the System Operator (or other privileged users) to:

- a. Abort the printing of a file
- b. Backup and reprint a page (for paper changes)

3.1.1.2.2.3 Assign/Deassign

The Assign/Deassign functions shall allow users to become owners of certain physical devices without having the privilege to do general ownership changing. Assign shall only allow the assignment of devices from a pool of free devices. It shall not be possible to assign a device that is presently in use or which is reserved for the system or another user.

Assign and Deassign shall be spawned through the Secure Server mechanism. The user's interface to Assign shall consist of specifying the device(s) to be assigned. The user interface to Deassign shall consist of specifying the device(s) to be deassigned. The System Operator shall be able to force a device to be deassigned.

The ability to assign ownership of a physical device such as a tape or disk drive to a user shall allow the user to use that device to access his media without risk of unauthorized access by other users. The user should assign a device before physically placing his media in it. This mechanism shall protect media from unauthorized access between the time the media is physically mounted and the time the user process opens the device.

3.1.1.2.2.4 Mount/Unmount

Mount and Unmount shall be spawned through the Secure Server mechanism. The user shall supply the name of the file system to be mounted, the drive on which the file system resides, the pathname of a presently accessible entry where the file system is to be mounted (normally, the file would be of the directory subtype), and an optional argument indicating that the file system is to be mounted read only.

The user shall have discretionary access to the extent (i.e., if the extent is to be mounted read/write, the user must have read/write access; if the extent is to be mounted read only, the user must have read access) and write discretionary access to the entry where the file system is to be mounted to be able to

Non-Kernel Security Related Software B5 Specification

successfully request that it be logically mounted.

The interface to the Unmount Process shall also be similar to that in the current UNIX system. The user shall supply the name of the file system to be unmounted.

3.1.1.2.2.5 System Startup/Shutdown

The user interface to the System Startup Process shall consist of supplying the name and location (i.e. extents) of the Security Kernel and Initial Process images to be loaded. This interface is provided by a small bootstrap program which is loaded using the ROM bootstrap or other mechanisms. The Initial Process shall then ask the user for the location of the root file system and for the time.

The System Shutdown Process shall be invoked by logging in as a special user whose default "shell" is a system shutdown program. This login shall be permitted only from the system console, and shall be protected by the usual password protection mechanisms.

3.1.1.2.2.6 Secure Mail

The user interface to the Secure Mail function shall consist of two programs which allow for the sending and receiving of mail. The basic interface shall be similar to the interface provided by the current UNIX mail program. However, the receive mail portion shall only deliver mail that exists at the security level of the Secure Mail process. The user interface shall allow the user to send a file or input from his terminal to one or more other users identified by their login identifiers. Mail may be sent by users running at any security level and will be protected at that level. The receive part of the Secure Mail function shall allow a user to read all mail sent to him that is at his current security level. To read mail at another security level, the user shall be required to change his security level.

3.1.1.2.3 System Maintenance Services

The System Maintenance functions shall be spawned through the Secure Server mechanism. Input parameters and output are defined in the Detailed Functional Requirements below.

3.1.1.2.4 System Administrative Services

The System Administrative Services class of Non-Kernel Security Related Software shall consist of the functions to support the administrative operation of the system, such as adding and deleting users.

3.1.1.2.4.1 User Control Editor

The User Control Editor shall be used by the System Administrator to control user access to the system. Specifically, the process shall be used by the System Administrator to add and delete users and groups from the system, to change user and group access level capabilities, to change a user's initial

Non-Kernel Security Related Software B5 Specification

login group, and to set the identification of the Group Administrator for a group. These changes shall be exercised by updating the User Access Authentication Data Base and the Group Access Authentication Data Base. Since all decisions about access in KSOS are made on the basis of information contained in these data bases, the interface to this function shall be via the secure terminal path; i.e., the function shall be spawned off by the Secure Server. Alteration of these data bases shall not affect currently logged in users until they request some service for which one of the data bases would be checked. It is recommended that the System Administrator request (or force) a logout of presently active users affected by changes to these data bases.

3.1.1.2.4.2 Immigration Officer Process

The function of the Immigration Officer shall be to check the acceptability of "foreign" file systems. The immigration officer process shall take the name of the file system to be immigrated. It will restore (see File System Restore CPC 3.2.2) the file system from tape and checks its acceptability. Only after the Immigration Officer has approved the file system shall it be acceptable for mounting.

3.1.1.2.4.3 Audit Capture Process

The Audit Capture Process shall receive, via either the signal or the inter-process communication mechanisms, raw audit information from the Kernel and NKSR processes. It shall write this information to an audit file. The process shall be able to select the types of information to be kept and may add additional information such as a time stamp to the information before writing. Multiple audit files shall be used to save the audit information. The process shall accept a command from the System Operator to switch audit files.

3.1.1.2.4.4 Privilege Control Process

The Privilege Control Process shall provide the Security Officer with the ability to grant and revoke special access privileges for process image files. Privileges shall be given to processes by changing elements of the type independent information for the file from which the process is invoked. The privileges of presently active processes shall not be changed by the Privilege Control Process. The following privileges shall be capable of being granted or revoked:

- a. The ability to violate the simple security property. (This privilege is not presently used. It is included for completeness.)
- b. The ability to violate the simple integrity property.
- c. The ability to violate the *-property for security.
- d. The ability to violate the *-property for integrity.
- e. The ability to violate discretionary access.
- f. The ability to use the K_secure_terminal_lock primitive.

Non-Kernel Security Related Software B5 Specification

- g. The ability to use the `K_release_openers` primitive.
- h. The ability to use the `K_link` and `K_unlink` primitives.
- i. The ability to use the `K_mount` and `K_unmount` primitives.
- j. The ability to change the following type independent information about an object:
 - 1. The security level of the object.
 - 2. The integrity level of the object.
 - 3. The owner of the object.
- k. The ability to set type dependent information (status) about an object; i.e., the ability to open an object in status update mode.
 - 1. The ability to save the process image in the swapping area; i.e., the ability to set the "sticky bit" (ISVTXT).
- m. The ability to lock a segment in core (not swappable).
- n. The ability to grant and revoke privileges.

3.2 Detailed Functional Requirements

3.2.1 Secure User Services

A set of functions shall be provided which communicate with the user terminal through a secure communications path. These functions shall be known as Secure User Services.

3.2.1.1 Secure Initiator

3.2.1.1.1 Users and Privileges

The Secure Initiator shall be invoked by the Kernel as part of the System by the Initial Process as part of system initialization.

3.2.1.1.2 Inputs

The Secure Initiator will receive, from the Initial Process, the seeds for the user and NKSR process bootstrapper segments.

3.2.1.1.3 Processing

The Secure Initiator shall spawn off the Audit Capture Process and shall make UNIX directory entries for the Process Bootstrapper segments. It shall then initialize the mount table and the access levels of various subtypes and devices. At this point, the Secure Initiator shall execute a shell file of commands similar to the rc files on UNIX. Finally, the Secure Initiator shall set the level of each configured terminal and spawn off a Secure Server for each

such terminal.

The Secure Initiator shall also set up the terminal data base which shall contain a list of terminals coupled with the Secure Server running on that terminal. This data base shall then be used by the Secure Servers.

3.2.1.1.4 Outputs

The Secure Initiator shall pass to each Secure Server the identification of the terminal on which it runs.

3.2.1.2 Secure Server

3.2.1.2.1 Users and Privileges

A separate Secure Server shall be spawned to service each terminal.

This Server shall be privileged to use the `K_secure_terminal_lock` Kernel call, to change certain type independent information about an object, and to change process type independent information.

3.2.1.2.2 Inputs

The process shall receive input from the Secure Initiator in the form of an argument segment, and from the user terminal via the secure terminal path.

Inputs from the user shall consist of the name of the service to be performed and an arguments which may be needed by that service.

3.2.1.2.3 Processing

The Secure Server shall prompt the user and read from the secure terminal path to determine type of service required. If the Server is unable to determine the type of service required, it shall display a rejection request on the user's terminal and thaw the terminal. In the case of operator requests, the Server shall determine that the user is currently running at OPERATOR integrity level or higher before an the function is performed. In the case of administrator requests, the Server shall determine that the user is currently running at ADMINISTRATOR integrity level before an the function is performed.

The Secure Server shall process user requests using a combination of spawned processes as needed to service user requests. Processing required for each type of service is described in the processing section for that service.

3.2.1.2.4 Outputs

The Secure Server shall communicate to the terminal user via the secure terminal path. Upon receiving an IPC signal from the Initiator the Secure Server shall display the user's current security level as a character string sent out on the secure path. Process which are spawned by the Secure Server shall receive their arguments via an argument segment.

Non-Kernel Security Related Software B5 Specification

Outputs required for each type of service are listed in the outputs section for that service.

3.2.1.3 Access Authentication Functions

The Access Authentication Functions shall process requests for login, change group, setting terminal access level, and logout. These functions shall be internal to the Secure Server process.

3.2.1.3.1 Inputs

Valid user names and identifiers, encrypted passwords, and the user access matrix shall be available from the User Access Authentication Data Base.

Valid group names and identifiers, encrypted passwords, and the group access matrix shall be available from the Group Access Authentication Data Base.

Profile information for each terminal shall be available from the Terminal Profile Data Base. The user's name, unencrypted password, and group name shall be obtained via the secure terminal path.

3.2.1.3.2 Processing

The Access Authentication Functions shall process user requests for logging in, logging out, changing group, and changing access level. Additionally, they shall process certain requests from the System Operator which control login/logout.

3.2.1.3.2.1 Login Processing

After receiving a login request the current state of the terminal shall be checked. If the terminal is currently logged-in, the request shall be denied. If the terminal is not logged-in a message shall be displayed on the terminal requesting the user's login name. After obtaining the user's login name, the process shall disable echoing on the terminal and request the user's password. The password entered by the user shall then be encrypted and the user's login name and encrypted password shall be checked against the User Access Authentication Data Base. If there is no match, the login name and password will be requested again. Only a small number of retries shall be allowed. If this retry limit is exceeded the development contractor shall provide a mechanism for the notification of the System Security Officer. If there is a match, the processing shall continue. If the User Access Authentication Data Base indicates a login group, the group shall be checked against the Group Access Authentication Data Base and if valid, shall be used as the user's group without requiring a group password. If no group is indicated, the system default group shall be used. All users shall be initially logged in at their default security and integrity level unless that level is above the current maximum permitted on the system, in which case the user's level shall be set to the system low, and he shall be informed of his level (system low). The following functions shall be performed in logging in a user:

- a. Select a terminal path, change the security and integrity levels for that path to the login level of the user, and the owner of the path to the user.

Non-Kernel Security Related Software B5 Specification

- b. K_spawn a process at the user's login level which will cause the user's shell to start running on the UNIX Emulator with the selected user path as the standard input and output. The use of the K_spawn primitive shall remove any special privileges possessed by the Secure Server.

3.2.1.3.2.2 Changing Group

After receiving a request to change group, a message shall be displayed on the user's terminal requesting the name of the group desired. The group name shall be checked against the Group Access Authentication Data Base to determine if the group is valid and if the user's present access level is permissible for the requested group. If a password for the group is listed the user shall be asked the password which shall be encrypted and checked. If successful, the requested group ID shall be set for the user, otherwise a rejection message shall be displayed on the user's terminal. As in login processing, only a limited number of attempts shall be permitted before notification of the System Security Officer.

3.2.1.3.2.3 Change Access Level

Upon requests by the user to change his current access level, the Server shall check that the requested level is valid for the user, group, and terminal by checking the appropriate data bases. The Server shall also verify that the requested level is permitted on the system at that time (i.e. that the requested level is at or below the current system high level). The Server shall then determine whether or not the user has an active environment (usually a shell running on a UNIX Emulator) at this level. If so, the Server shall select the user path with the previous environment. If the user does not have an active environment at this level, the Server shall attempt to find an inactive user path for that terminal and if successful, set the path level to the desired level and spawn a shell and emulator at that level.

3.2.1.3.2.4 Logout

Requests for logout shall cause a signal to be sent to all user processes invoked for that terminal requesting them to terminate. User processes which do not terminate within a reasonable time will be killed. The terminal mode (speed, parity, etc) shall be set to the values listed for it in the Terminal Profile Data Base.

3.2.1.3.2.5 Processing Operator Requests

The Access Authentication Functions shall accept and enforce commands from the System Operator to control user login, access levels, and logout.

3.2.1.3.3 Outputs

An accounting message shall be sent to the Audit Capture Process for each request even if the request was denied.

Non-Kernel Security Related Software B5 Specification

3.2.1.4 Password Modification Function

3.2.1.4.1 Inputs

The Password Modification Function shall have read/write access to the User Access Authentication Data Base and the Group Access Authentication Data Base. The Password Modification Function shall communicate with the user via the secure terminal path.

3.2.1.4.2 Processing

After receiving a request for password modification, the Secure Server shall invoke the Password Modification Function to handle the request. The Password Modifier shall turn off echoing on the user's terminal and request the user's old password which is then encrypted and checked against the User Access Authentication Database. This check is done to insure that the person on the terminal is the real user for the logged in user ID. The user will then be requested to enter the desired new password. The user shall be requested to enter the new password a second time. If the two entries do not match the user shall be informed and the process repeated. The new password shall be encrypted and written into the User Access Authentication Data Base. The user shall be allowed to modify a group password in a similar manner if he has been designated in the Group Access Authentication Database as the Group Administrator for that group.

3.2.1.4.3 Outputs

The Password Modification Function shall write to the user terminal via the secure terminal path.

The Password Modification Function shall write new passwords in the User Access Authentication Data Base, and the Group Access Authentication Database as required.

Audit information noting that the password has been changed will be captured.

3.2.1.5 File Access Modification Function

User requests to change the access level and owner of a file shall be processed by the File Access Modification Function.

3.2.1.5.1 Inputs

The File Access Modification Function shall receive input from the user via the secure terminal path including the pathname of the file to be modified.

The File Access Modifier shall have access to status information on all files.

3.2.1.5.2 Processing

The File Access Modifier shall request from the user the path name of the file to be modified. It then checks for the existence and ownership of the file. If the requester can not access the file, or the requester is not the owner, or the requester, if not the owner, is not at ADMINISTRATOR level, or the file is open, the request shall be denied. If these checks all succeed, the user shall then be asked the operation to be performed.

For a request to change access level, checks shall be made to determine that the file exists, that the requested level is allowed on the file system on which the file resides and that the requester's access matrix includes the requested level. If these checks are successful the level of the file shall be changed.

For changing the owner of a file, the Server shall check that the user is the owner of the file or is at the ADMINISTRATOR or higher integrity level, that the file exists, that the file is not open, and that the new owner exists. If these checks are successful, the owner of the file shall be changed.

3.2.1.5.3 Outputs

The File Access Modifier shall display the current level, discretionary access information, and owner prior to changing any of this information. In addition, the proposed new values shall be displayed prior to actually changing them. In the case of a change owner a message shall be sent to the new owner informing him that he has become the file owner. All output to the user terminal shall be via the secure terminal path.

Audit information shall be gathered on each request even if the request was denied.

3.2.2 System Operation Services

3.2.2.1 Network Daemon

The computer network interface functions are spread across the UNIX Emulator, the Security Kernel and the Network Daemon. Table 1 shows the overall functional division. The remainder of this paragraph discusses the functions performed by the Network Daemon.

Non-Kernel Security Related Software B5 Specification

Table I
TCP Functional Allocation

	Security Kernel	UNIX Emulator	Network Daemon
IMP Support	x		
Rendezvous			x
Flow Control		x	
Data Stream Integrity:			
Checksum			x
Resequencing		x	
Retransmission		x	
Acknowledge		x	
Duplicate Detect		x	
Routing			x
Connection Alloc.			x
Urgent Processing		x	
Formatting		x	
Security Level of Message			x

3.2.2.1.1 Users and Privileges

The Network Daemon shall be a general service program utilized by all users of the network. The Network Daemon shall be privileged to violate the security *-property to allow it to service a multi-level network. If the network is at only one security level, the Network Daemon is unprivileged.

3.2.2.1.2 Inputs

The inputs to the Network Daemon are the data streams to and from the network, and the control requests from the System Operator.

3.2.2.1.3 Processing

This paragraph discusses the functions of Table I which are performed by the Network Daemon.

Non-Kernel Security Related Software B5 Specification

3.2.2.1.3.1 Rendezvous

Rendezvous is the process of recognizing the match of connection names. For example, a remote user supplies, as part of his connection attempt, a local TCP port name to which he wishes to be connected. The process of matching the local name (on which some local process is listening) to the name being supplied by the remote process is called rendezvous. To accomplish the rendezvous, the Network Daemon shall maintain a table of connections and their names, both local and remote.

The Network Daemon shall provide the management of the internal tables necessary to provide the overall network services. These tables shall include, but are not limited to, the host status tables, the table containing the connection status for each open connection, and tables containing internal status information.

3.2.2.1.3.2 Checksum

To allow for proper internal routing, the Network Daemon shall compute the checksum of incoming messages. Messages received with an incorrect checksum shall be discarded in keeping with the TCP definition [Postel 78a] [Postel 78b]. The Network Daemon shall also compute the checksum for all outbound messages. The formula to be used for checksum calculation will be supplied by the Government.

3.2.2.1.3.3 Routing

The Network Daemon shall supply the routing information required by the network by mapping the user specified destination into the appropriate coded form.

3.2.2.1.3.4 Connection Allocation

The Network Daemon shall control the allocation of the connection resource. This allocation shall include the maintenance of internal tables describing the state of each connection. For active connections, this state also includes the information necessary to route inbound packets to the appropriate process and to fill in the fields for the destination on outbound traffic. The System Operator may manipulate the state of connections via the System Operator interface. This latter mode of operation is expected to be used infrequently.

3.2.2.1.3.5 Message Security Level

The Network Daemon shall supply the security level field of outbound messages based on the Kernel-supplied value for the level of the sending process if this field is supported by the network. If the network does not support messages of differing classification (e.g., a single level network) the Network Daemon shall not transmit messages whose security level is above the level of the network. The Network Daemon shall examine the security level of incoming messages to assure that they can be delivered to the receiving process without violation of the mandatory security policy. The security level of a message may be implicit, as in the case of single level networks.

Non-Kernel Security Related Software B5 Specification

In the case of single level networks, the Security Kernel shall provide many of the message security level functions by permitting (and denying) inter-process communication attempts to the Network Daemon. In these cases, the Network Daemon may be a process running at the level of the network. Attempts by processes of higher security level to communicate with the Network Daemon would be denied by the Security Kernel.

3.2.2.1.4 Outputs

The outputs of the Network Daemon are the data streams to and from the network and any diagnostic messages.

3.2.2.2 Line Printer Spooling

3.2.2.2.1 Users and Privileges

The Line Printer Spooler shall be an unprivileged service function available to all system users.

The Line Printer Daemon (de-spooler) shall be responsible for the printing of the appropriate classification markings on its output and shall be privileged to violate the security *-property to remove printed spool files.


3.2.2.2.2 Inputs

The inputs to the Line Printer Spooler shall be the files to be printed (which defaults to the standard input if not supplied) and the control information discussed in section 3.1.1.2.2.2. The inputs to the Line Printer Daemon shall be the previously spooled output files and control commands from the System Operator.

3.2.2.2.3 Processing

The overall flow of processing in the Line Printer Spooler shall be as shown in Figure 3-2.

Non-Kernel Security Related Software B5 Specification



* set uid on execution

Figure 1-3

Non-Kernel Security Related Software B5 Specification

The spooler shall make a copy of the files to be spooled, which, together with any control information, shall be placed in the spool directory. This directory, as well as the files it contains, shall have a discretionary access protection which disallows access to all users other than "spool", the spooler manager.

Using the corresponding control information, the Line Printer Daemon shall print each spooled file. The Daemon shall apply the appropriate headers, footers, and burst page information to correctly mark the classification of the output. Markings shall be in accordance with DoD Directive 5200.1-R. The Daemon shall create auditing records noting the creation of classified output. Finally, the Daemon shall delete the spooled output file after the printing has been successfully completed.

The Daemon shall respond to inter-process communication from the System Operator, as defined by the Daemon Control Interface above. The Daemon shall be capable of aborting the printing of a file and of backing up a page for paper changes.

3.2.2.2.4 Outputs

The output of the Spooler shall be the temporary files to be subsequently printed by the Daemon. The output of the Daemon shall be the printed output and the auditing records indicating the creation of classified output.

3.2.2.3 Assign/Deassign

3.2.2.3.1 Users and Privileges

The use of assign shall be unrestricted although the ability to assign certain devices may be restricted. The Assign and Deassign functions shall be privileged to change the owner and level of a device; i.e., to set the type independent information for the device.

3.2.2.3.2 Inputs

The user input to Assign shall specify the device(s) to be assigned. The input to Deassign shall indicate the device(s) to be deassigned. Both processes shall have access to the type independent status of devices.

3.2.2.3.3 Processing

Upon receipt of a request to assign a device, the Assign process shall check the current owner of the device. If the device is not owned by the System Resource Manager, the process will assume the device is presently assigned and disallow the request since a device can only be assigned to one user at a time. If the level of the user requesting the device is higher than the maximum level allowed for that device, the request will also be denied. To assign a device, Assign shall change the owner and access level of the device to that of the requester and the discretionary access permissions to read/write by owner.

Deassignment of a device must be requested either by the current device owner or the System Operator. The device ownership shall be return to the System

Non-Kernel Security Related Software B5 Specification

Resource Manager upon deassignment.

3.2.2.3.4 Outputs

The Assign and Deassign processes shall output messages as necessary to the user.

3.2.2.4 Mount/Unmount

3.2.2.4.1 Users and Privileges

The use of the Mount and Unmount Processes shall be unrestricted, however, the ability to mount and unmount particular file systems or to mount on particular devices may be restricted.

The Mount Process shall have the permissions needed to access the underlying devices.

3.2.2.4.2 Inputs

The user inputs to Mount shall be the name of the file system to be mounted, the drive on which the file system resides, the pathname of a presently accessible file (normally, the file would be of directory subtype), and an optional argument indicating that the file system is to be mounted read only. If no arguments are given, a list of the currently mounted file systems is given.

The process shall also read a file, maintained by the Immigration Officer Process, containing the file systems which are acceptable for mounting.

The Unmount Process inputs from the user shall be the drive to be unmounted.

3.2.2.4.3 Processing

The Mount Process shall verify that the file system is acceptable for mounting by checking that the file system is on the list maintained by the Immigration Officer Process and that it is not already mounted. Mount shall also check that the user has discretionary access to the file system (i.e., if the file system is to be mounted read/write, the user must have read/write access; if the file system is to be mounted read only, the user must have read access), that the user has write access to the parent directory, that the user's security and integrity levels are greater than or equal to the file system's level, and that the file system's set of security compartments is a subset of the user's set. If any of these checks fail, the process shall send a message to the Audit Capture Process and notify the user of the failure. If the checks are successful, the Mount Process shall issue the K_mount Kernel call which will logically mount the file system. The UNIX directory manager (UDM) ZZZ will then be called to add the file system into the UNIX structure.

The Unmount Process shall check that the user has permission to unmount the file system; i.e., the user has discretionary access to the file system, has write access to the mount on entry of the directory being unmounted, and has the

Non-Kernel Security Related Software B5 Specification

appropriate security level, integrity level, and security compartments as described in the section on Mount. If these checks are successful, the Unmount Process shall issue the K_unmount Kernel call. Followed by a call to the UNIX directory manager to remove the file system from the UNIX structure.

3.2.2.4.4 Outputs

The Mount and Unmount Processes shall update a file analogous to the file /etc/mount in current UNIX which is used to inform users which file systems are presently mounted. Both processes shall provide diagnostic messages as required.

3.2.2.5 System Startup/Shutdown

3.2.2.5.1 Users and Privileges

The use of the System Startup and Shutdown mechanisms shall be restricted to the System Operator. Since the KSOS system is not in operation when the System Startup mechanism is invoked, and, thus, cannot protect itself, the System Startup mechanism shall be further restricted to be used only from the operator's console. The System Shutdown mechanism shall be similarly restricted because the KSOS system will be unavailable to protect itself after the shutdown is complete.

3.2.2.5.2 Inputs

The inputs to the System Startup Process shall be the name of the KSOS image and of the Initial Process image to be loaded. Locally specified functions may require additional inputs. The Initial Process shall then ask the user for the time, the location of the root file system, and the Secure Initiator to be used.

No inputs to the System Shutdown Process are required. Locally specified functions may require additional inputs.

3.2.2.5.3 Processing

The System Startup Process shall begin by having the ROM bootloader (or equivalent) load a KSOS bootstrap. This bootstrap shall then load the Kernel and Initial Process, and initialize the system hardware to allow proper execution of the Kernel. The Kernel will perform any internal initialization required. It will then set up internal tables, etc. to allow the Initial Process image to run as a user process, and will begin running it.

The Initial Process shall set the time of day clock, mount the root file system, and build the segments used by the Process Bootstrapper (PBB). The process shall then issue a K_invoke call to invoke the Secure Initiator.

The System Shutdown Process shall first execute a locally modifiable shell file to accomplish any desired local processing. It shall then "kill" all user processes, force all cached I/O to be completed, and call K_halt to halt the kernel.

Non-Kernel Security Related Software B5 Specification

3.2.2.5.4 Outputs

The System Startup and Shutdown Processes shall have no outputs other than diagnostic and informational messages. Sites may elect to include additional output from their locally specified startup and shutdown shell files.

3.2.2.6 Secure Mail

3.2.2.6.1 Users and Privileges

The Secure Mail mechanism shall be a service available to all users. It shall be unprivileged.

3.2.2.6.2 Inputs

The input for the send part of Secure Mail shall be the name of a file to be mailed (which shall default to the standard input if not supplied), and the list of users to whom the mail is to be sent.

The input for the receive part of Secure Mail shall be the previously transmitted letters, if any, and an (optional) file name in which to save the mail.

3.2.2.6.3 Processing

Conceptually, Secure Mail shall be similar to "two ended spooling". The sending process shall place the mail in a "post office box" directory that is unique for each user. This directory as well as the files (letters) placed in it, shall be owned by the user "post" and shall have discretionary access protections that disallow access to all users other than "post".

The receive mail process shall retrieve the letters from the user's "post office directory" and place them in a user specified directory. So that security upgrading of letters does not occur, the user shall only be allowed to receive letters as his current security level. The mail may then be examined by the user.

3.2.2.6.4 Outputs

The output of the send part of Secure Mail shall be the letters created under the /post/to_user directories, and any diagnostic information required, such as that the specified recipient does not exist. The output of the receive part shall be a file in the user's mail directory.

3.2.3 System Maintenance Services

This class of Non-Kernel Security Related Software shall consist of those functions that provide facility for:

- a. file system backup and restoration
- b. file system consistency check and repair

Non-Kernel Security Related Software B5 Specification

c. system distribution, generation, and reconfiguration

3.2.3.1 File-system Backup and Restoration

The Dump and Restore functions shall provide a mechanism for the generation and recovery of incremental and complete file system dumps of KSOS files. Both Dump and Restore programs shall be trusted to allow the reading and creation of multi-level file objects.

3.2.3.1.1 Dump

3.2.3.1.1.1 Users and Privileges

Use of the Dump program shall be restricted to the Operator level and above. The program shall have the permissions needed to access the disk devices directly (rather than going through the Security Kernel and UNIX Emulator file systems).

3.2.3.1.1.2 Inputs

The inputs to the Dump program shall be similar to the inputs to the current UNIX dump program:

- a. parameters which describe precisely what is to be dumped
- b. the name of the file system(s) to be dumped
- c. the data from the file system being dumped

3.2.3.1.1.3 Processing

After determining that the invoking user has a sufficient access level to run the Dump program, the program shall copy a KSOS file system to its output in a dump format. Dump shall be capable of producing an incremental dump of a specified KSOS file system. The incremental dump shall include all files that have changed since the last full dump, or in a specified period.

3.2.3.1.1.4 Outputs

Output of the Dump program shall be to a 9-track magnetic tape or other removable storage medium. The Dump program shall allow its output to be directed to a disk or equivalent.

3.2.3.1.2 Restore

3.2.3.1.2.1 Users and Privileges

Use of the Restore program shall be restricted to the Operator level and above. The program shall have the permissions needed to access the disk devices directly (rather than going through the UNIX Emulator file systems).

3.2.3.1.2.2 Inputs

The inputs to the Restore program shall be similar to those used by the current restor(VII) [Thompson 75] program:

- a. parameters describing precisely what is to be restored
- b. the name of the file system to be restored
- c. the data previously dumped that is to be restored

3.2.3.1.2.3 Processing

After validating that the user is privileged to restore a file system, the Restore program shall (selectively) restore to the designated file system.

3.2.3.1.2.4 Outputs

The output of the Restore program shall consist of the restored files or file system, and any diagnostic messages produced.

3.2.3.2 File System Maintenance Services

3.2.3.2.1 Make File System

3.2.3.2.1.1 Users and Privileges

Operation of the Make File System program shall be allowed only by the System Administrator since the contents of the target pack, extent, or file system are destroyed by the initialization process. The program shall have the permissions needed to access the underlying device files, rather than going through the Security Kernel's or UNIX Emulator's file systems.

3.2.3.2.1.2 Inputs

The inputs to the Make File System program shall include the parameters which describe the initial state of the pack, extent, or file system which is being initialized.

3.2.3.2.1.3 Processing

The Make File System program shall initialize the pack, extent, or file system in accordance with the parameters supplied. The program may implement safety measures to reduce the probability of accidentally re-initializing a file system.

3.2.3.2.1.4 Outputs

The outputs of the Make File System program shall include the newly initialized pack, extent, or file system, and any diagnostic or warning messages.

Non-Kernel Security Related Software B5 Specification

3.2.3.2.2 Directory Consistency Check

3.2.3.2.2.1 Users and Privileges

The use of the Directory Consistency Check program shall be limited to the System Operator level and above since it requires access to the underlying device(s) on which the file system is built.

3.2.3.2.2.2 Inputs

The input to the Directory Consistency Check program shall include the name of the file system to be checked.

3.2.3.2.2.3 Processing

The program shall read the UNIX file system directories on the device, and compute the actual and expected link counts for each file. Counts that do not match shall be reported. In addition, any file which has an actual and an expected count of zero shall be reported.

3.2.3.2.2.4 Outputs

The outputs of the Directory Consistency Check program shall include the diagnostic messages discussed above. The program shall also report the absence of anomalies in the structure of the file system.

3.2.3.2.3 Storage Consistency Check

3.2.3.2.3.1 Users and Privileges

The use of the Storage Consistency Check program shall be limited to the System Operator level and above since it requires access to the underlying device(s) on which the file system is built.

3.2.3.2.3.2 Inputs

The inputs to the Storage Consistency Check program shall include the device to be checked and parameters to control the optional processing.

3.2.3.2.3.3 Processing

The program shall examine the named file system to determine that all allocated areas of the file system are mutually exclusive and that all storage areas are accounted for in the file system. The program shall have the capability to identify files that include named blocks.

3.2.3.2.3.4 Outputs

The output of the Storage Consistency Check program shall include a summary of the storage allocation on the device including any diagnostic information about missing or multiply allocated areas.

3.2.3.2.4 Kernel Name to Pathname Mapping

3.2.3.2.4.1 Users and Privileges

The use of the Name Mapping program shall be restricted to the System Operator because it must access the underlying device on which the file system is built.

3.2.3.2.4.2 Inputs

The inputs to the Name Mapping program shall be a list (possibly null) of Secure Entity Identifiers (SEID's) (the Security Kernel CPCI's names for objects), and the device (file system) on which they are to be found.

3.2.3.2.4.3 Processing

If no Secure Entity Identifiers have been supplied the program shall produce a list the UNIX path names vs. SEID's for all files on the device. Otherwise, the list shall include only the requested SEID's

3.2.3.2.4.4 Outputs

The outputs of the Name Mapping program shall be the list of UNIX path names vs. SEID's, and diagnostic information as required.

3.2.3.2.5 Modify Control Entries for Filesystem

3.2.3.2.5.1 Users and Privileges

The use of the program which can modify the control entries for a file (MCE) shall be restricted to the System Operator since it requires access to the device on which the file system resides. In addition, the program should only be used by personnel with sufficient knowledge of the file system structure to be aware of exactly what may and may not be safely changed.

3.2.3.2.5.2 Inputs

The inputs to MCE shall include the file system to be modified, the node to be modified, and the fields to be changed.

3.2.3.2.5.3 Processing

The MCE program shall modify the requested fields to the requested value. The program should produce warning messages about how serious the proposed modification may be.

3.2.3.2.5.4 Outputs

The output of MCE shall be the modified file system. In addition, the file system shall be removed from the list of acceptable systems for mounting, forcing the Immigration Officer Process (3.1.4.4.4) to be invoked before the file system may be mounted. Finally, the program shall produce audit records of its use.

Non-Kernel Security Related Software B5 Specification

3.2.3.3 KSOS System Generation

Each KSOS installation can be expected to have a unique hardware configuration. Due to the requirement that KSOS must support a variety of devices, it is not practical to distribute a single "universal" copy of the system. Rather, each site will generate a system tailored to its hardware configuration.

3.2.3.3.1 Users and Privileges

The System Generation mechanism shall be designed to be used by an individual with minimal computer system expertise (e.g., the Security Officer who must certify the resulting system). The function shall possess no special privileges. However, it must be recognized that access to the KSOS system image as it is being built is tantamount to being able to violate all the rules of the system.

3.2.3.3.2 Inputs

The KSOS system shall be distributed as a three-component package. The first component, a distribution tape, shall contain a starter system and standard system-build libraries to be copied to the target system disk at the remote site. Documentation for KSOS shall comprise the second component of the distribution. The third element of the distribution package shall consist of the system generation instructions and generation password sequence. The distribution tape and generation package shall be transmitted separately as SECRET-level material. The distribution tape will normally be sent to the site manager, while the installation guide and passwords will be sent to the System Security Officer.

A checksum technique shall be used to validate transfer of the Kernel components during the system build cycle. Each copy of the Kernel modules shall have a unique identifier (probably derived from the target machine's serial number) that will generate unique module checksums for each distribution tape. Each element in the distributed system shall contain a component checksum. An interactive system-build program shall compute the checksum associated with each component, notifying the system builder (Security Officer) of the checksum failure and identifying the failed component. As a general rule, systems built from "failed" modules cannot be certified as multi-level secure. Using these concepts, the required integrity can be provided to remotely generated systems without requiring the System Security Officer to be a system programmer.

The standard KSOS distribution tape shall contain a bootstrap loader, a site-specific starter system used to generate the target KSOS system and all modules required to generate a KSOS system, at a remote site.

3.2.3.3.3 Processing

3.2.3.3.3.1 Bootstrap

The bootstrap loader shall be the first entry on the KSOS distribution tape. A small program entered from the machine console (or executed from a standard ROM loader) shall read the bootstrap which shall, in turn, read the starter system from the tape, and then transfer control to the starter system. Each of the tape-based programs shall be checksummed to detect media failure.

Non-Kernel Security Related Software B5 Specification

3.2.3.3.2 Starter System

The starter system shall be a KSOS system previously generated for a particular minimum configuration. The starter system shall require no more than the minimum hardware specified in the System Specification (Type A). The starter system may be specific to a particular hardware configuration. When the distribution tape is generated by the maintenance and support organization, the appropriate starter system shall be selected. Alternatively, the distribution tape may contain starter systems for several common configurations. The device and vector addresses of the disk drive and console shall have standard values. Any other configuration shall be handled as a special case by the maintenance and support organization.

3.2.3.3.3 Libraries

The starter system shall transfer the KSOS system programs, libraries, and archives from the distribution tape to the specified system disk to complete the system restoration process. Media failures should be detected through component checksums.

3.2.3.3.4 Special Cases

Unique configurations with no magnetic tape or non-standard system disks or console addresses shall be handled as special cases. Installations without tape drives shall receive a disk-based distribution of KSOS.

3.2.3.3.5 System Generation

System generation is basically a matter of selectively including the appropriate device drivers in the Security Kernel, defining the /dev files and their underlying Kernel representations and "binding" the /dev files to the Kernel's assignment structure. The remainder of the system software may require minor local (per site) modifications, for example to set the local time zone. All such common local modifications shall be automated through the use of shell scripts or other equivalent mechanisms.

The Security Officer will execute an interactive System Configurator program similar to mkconf(VIII) [Thompson 75] for standard UNIX. The System Configurator program may be executed under the starter system or under a larger KSOS configuration. The Security Officer shall interact with the program to define the hardware configuration for which the KSOS system is to be generated. A list of hardware parameters (e.g., in the PDP-11/70 version, base address of the control and status registers, vector addresses, and interrupt priorities) must be provided to the Security Officer by a knowledgeable individual. A "standard" list shall be provided in the installation instructions.

The System Configuration dialogue shall allow:

- a. selection of devices.
- b. specification of device control registers, and vector addresses, and interrupt priority levels.

Non-Kernel Security Related Software B5 Specification

c. device driver source/entry names.

The System Configurator program shall prompt with "standard" device information, unless a local configuration file is specified. In that case, the Configurator program shall be used to create an updated (new) configuration file. The new configuration file shall be checked for consistency, and then used to generate the "new" local system by comparison of the new and old configurations. Any changes shall cause the appropriate drivers to be edited, recompiled, checksummed, and archived. Appropriate controls on use of the Configurator program must be observed, since portions of the Security Kernel will have to be recompiled to change device register addresses.

3.2.3.3.3.6 Integrity of the Inputs

The configuration procedure shall execute the internal program that validates the checksum of each of the modules required in the link edit phase. In the event of a checksum error, the failed modules shall be identified to the system builder (Security Officer). If any checksum failures are noted, the system shall not be certified as being suitable for multi-level secure operation.

3.2.3.3.3.7 Rename System

The final operation of the configuration procedure shall give the system its bootstrap name, e.g. KSOS-1.0. This step should not occur if any errors (e.g., checksum errors) have occurred. This procedure is executed prior to the Security Officer bootstrapping up the new version of the KSOS system. For additional security and integrity it is recommended that sites do not leave the sources for the trusted software on normally on-line disks.

3.2.3.3.3.8 System Regeneration

When new hardware is added to a configuration, the system will need to be "regenerated" to reflect the changes. This regeneration process should be carried out by the Security Officer in a single-user environment. By using the Configurator program in an update mode, only new information or changes will need to be entered to create the new system configuration. While the procedure could easily be carried out in a multi-user environment, single-user operation appears to be an appropriate mode of operation during system regeneration of KSOS.

3.2.3.3.4 Outputs

The output of the system (re)generation process shall be a "bootstrappable" KSOS image, and any diagnostic or status information required.

3.2.4 System Administration Services

3.2.4.1 User Control Editor

The User Control Editor shall be used by the System Administrator to control access to the system by modifying the User Access Authentication Data Base and the Group Access Authentication Data Base.

Non-Kernel Security Related Software B5 Specification

3.2.4.1.1 Users and Privileges

Use of this program shall be restricted to users running at the system high security and integrity level and in the system compartment.

This process shall be privileged to issue `K_secure_terminal_lock` calls.

3.2.4.1.2 Inputs

The User Control Editor shall receive input from the Administrator via the secure path. It shall have read/write access to the User Access Authentication Data Base and to the Group Access Authentication Data Base. The User Control Editor shall also have read access to the Security Map Data Base.

3.2.4.1.3 Processing

The User Control Editor shall process requests from the System Administrator to add and delete users, to change a user's default login level, to change a user's access level capabilities, to change a user's initial login group, to change user and group passwords, and to assign Group Administrators. Appropriate checks shall be made by the process to assure data base integrity. Audit messages shall be generated and sent to the Audit Capture Process.

3.2.4.1.4 Outputs

The User Control Editor shall send such messages as necessary to the user terminal via the secure terminal path. It shall also make changes as necessary to the User Access Authentication Data Base and the Group Access Authentication Data Base. In addition, it shall generate and send to the Audit Capture Process necessary audit trail messages.

3.2.4.2 Audit Capture Process

The Audit Capture Process shall be a system process which receives audit messages and writes them onto a accounting file.

3.2.4.2.1 Users and Privileges

The Audit Capture Process shall be a system process having no direct users.

3.2.4.2.2 Inputs

The Audit Capture Process shall receive accounting information from the Kernel and NKSR processes via the inter-process communication mechanism and the signal mechanism.,

The Audit Capture Process shall accept messages sent through the signal mechanism from processes running at OPERATOR integrity level or higher to execute certain System Operator commands.

Non-Kernel Security Related Software B5 Specification

3.2.4.2.3 Processing

All messages received by the Audit Capture Process shall be marked as to whether they were sent using the signal mechanism or the inter-process communications mechanism. The message type shall then be checked to determine if the type is valid. The information may be reformatted and additional information such as a time stamp added before the message is written on the accounting file.

The Audit Capture Process shall accept requests from the System Operator to switch accounting files.

3.2.4.2.4 Outputs

After processing, accounting records shall be written on the audit file.

3.2.4.3 Privilege Control Process

The Privilege Control Process shall provide the Security Officer with the ability to to grant and revoke special privileges to processes.

3.2.4.3.1 Users and Privileges

The Privilege Control Process shall be a restricted use process run by the System Administrator at system high security and integrity level and in the system compartment.

The Privilege Control Process shall have the ability to violate the *-property for security and integrity, and discretionary access.

3.2.4.3.2 Inputs

The Privilege Control Process shall receive input from the Security Officer via the terminal. It shall also be able to obtain status information on files.

3.2.4.3.3 Processing

The Privilege Control Process shall receive requests from the System Security Officer to grant and revoke special privileges to processes by changing elements of the object type dependent information for the file from which the process is invoked.

3.2.4.3.4 Outputs

The process shall write to the Security Officer via the secure terminal path. It shall change file status information as necessary. The Privilege Control Process shall also send audit messages as necessary.

3.2.4.4 Immigration Officer Process

The Immigration Officer Process shall maintain a data base of file systems may which be mounted on the system.

Non-Kernel Security Related Software B5 Specification

3.2.4.4.1 Users and Privileges

The Immigration Officer Process shall be a restricted use process and can be run only by a user at the ADMINISTRATOR integrity level.

The Immigration Officer Process shall have the permissions needed to access the underlying disk devices.

3.2.4.4.2 Inputs

The inputs to the Immigration Officer Process shall consist of the current version of the mountable file systems data base, parameters which specify which file system is to be immigrated, and the data making up the file system.

3.2.4.4.3 Processing

The Immigration Officer Process shall check that no privileged software resides on the volume. The development contractor may include additional checks. If all checks are successful, the program shall perform any remapping of security level representations that has been requested. The Immigration Officer shall first check that the proposed remapping is valid; e.g., that this system includes all of the access levels represented on the foreign file system. If all the checks and the remapping (if required) are successful, the immigration officer data base shall be updated to reflect the fact that the volume is acceptable for mounting.

3.2.4.4.4 Outputs

The outputs of the Immigration Officer Process shall be the updated immigration officer data base and any diagnostic messages required.

3.2.5 Special Requirements

3.2.5.1 Human Performance

This paragraph is not applicable to this specification.

3.2.5.2 Government Furnished Property List

The compiler used for the Non-Kernel Security Related Software may be furnished by the Government.

3.3 Adaptation

The Non-Kernel Security Related Software should require no per-site modification. Individual sites may, however, modify the non-trusted functions as they see fit. Field maintenance of the trusted functions shall be prohibited unless subject to the quality assurance procedures (e.g. formal specification, verification, etc) equivalent to its initial creation.

Non-Kernel Security Related Software B5 Specification

4. QUALITY ASSURANCE PROVISIONS

4.1 Introduction

The underlying philosophy of KSOS quality assurance is to provide a convincing demonstration of the security and completeness of the system. Testing and quality assurance are a complement to the formal verification aspects of KSOS. All testing on KSOS shall be conducted at the development contractor's facility.

KSOS shall be subject to the technical reviews and audit procedures of MIL-STD-1521A, Appendices A-F. The following technical reviews and audits will be conducted:

- a. System Requirements Review (Appendix A)
- b. System Design Review (Appendix B)
- c. Preliminary Design Review (Appendix C)
- d. Critical Design Review (Appendix D)
- e. Functional Configuration Audit (Appendix E)
- f. Physical Configuration Audit (Appendix F)

The Critical Design Review shall include a review of any mathematical proofs showing that the design meets the requirements of the Government approved DoD security model.

The Functional Configuration Audit shall be the vehicle for the formal review of the design versus the mathematical description. The Physical Configuration Audit shall be the vehicle for the formal review of the "as built" computer programs versus both their design and their supporting documentation. In addition, the Physical Configuration Audit shall verify that the implementation requirements of the System Specifications (Type A) have been met. Both audits shall include review of any relevant mathematical proofs of correctness.

4.1.1 Category I Testing

The development contractor shall establish test plans to demonstrate that all the requirements of Section 3 have been met. Because this CPCI is composed of a number of autonomous functions, the primary vehicle for this demonstration shall be the Preliminary Qualification Tests. These tests shall be designed to exercise all the interface options for each function. To the extent feasible, the tests shall exercise all combinations of options. The tests shall include sequences intended to fail, such as attempts to violate the system's security rules or to use unacceptable combinations of options. Government approval of all test plans is required.

The test sequences for trusted functions shall be designed with the same rigor as those for the Security Kernel. Specifically, the test sequence should attempt to assure that all statements of the program have been exercised.

Non-Kernel Security Related Software B5 Specification

The general mechanism for assuring that the requirements of Section 3 have been met shall be a comparison of the system "state" before and after the function has been invoked. To ensure repeatability, all testing shall be automated if possible. Such automation may include (but is not limited to) shell scripts, test programs, and terminal emulation by another computer system.

4.1.2 Category II Testing

The KSOS system shall be subject to an overall system test (Category II) in accordance with the System Specification (Type A) and the development contract. This Category II test shall demonstrate that the three CPCI's which make up the KSOS system function correctly together. The Category II testing shall attempt to represent a typical user load including at least ten (10) on-line terminals and a network connection. To the extent feasible, the Category II testing shall also be automated through the use of shell scripts, and (optionally) another computer system emulating the terminal load.

Non-Kernel Security Related Software B5 Specification

Table II - Quality Assurance

Assurance techniques:

NA	Not Applicable
V	Verification
F	Formal Specification
T	Testing
D	Demonstration
A	Analysis
I	Inspection

Applicable test types:

1	Module level testing
2	Integration testing
3	Acceptance testing

Requirement	Assurance Techniques							Test types		
	NA	V	F	T	D	A	I	1	2	3
3.1.1.1.				x	x	x	x	x	x	x
3.1.1.2.1.1.		x	x	x	x	x	x	x	x	x
3.1.1.2.1.2.		x	x	x	x	x	x	x	x	x
3.1.1.2.2.1.1.		x	x	x	x	x	x	x	x	x
3.1.1.2.2.1.2.				x	x	x	x	x	x	x
3.1.1.2.2.1.3.				x	x	x	x	x	x	x
3.1.1.2.2.2.		x	x	x	x	x	x	x	x	x
3.1.1.2.2.3.		x	x	x	x	x	x	x	x	x
3.1.1.2.2.4.		x	x	x	x	x	x	x	x	x
3.1.1.2.2.5.				x	x	x	x	x	x	x
3.1.1.2.2.6.				x	x	x	x	x	x	x
3.1.1.2.4.1.				x	x	x	x	x	x	x
3.1.1.2.4.2.				x	x	x	x	x	x	x
3.1.1.2.4.3.				x	x	x	x	x	x	x
3.1.1.2.4.4.				x	x	x	x	x	x	x
3.2.1.1.		x	x	x	x	x	x	x	x	x
3.2.1.2.		x	x	x	x	x	x	x	x	x
3.2.2.1.		x	x	x	x	x	x	x	x	x
3.2.2.2.		x	x	x	x	x	x	x	x	x
3.2.2.3.		x	x	x	x	x	x	x	x	x
3.2.2.4.		x	x	x	x	x	x	x	x	x
3.2.2.5.				x	x	x	x	x	x	x
3.2.2.6.				x	x	x	x	x	x	x
3.2.3.1.1.				x	x	x	x	x	x	x
3.2.3.1.2.				x	x	x	x	x	x	x
3.2.3.2.1.				x	x	x	x	x	x	x
3.2.3.2.2.				x	x	x	x	x	x	x
3.2.3.2.3.				x	x	x	x	x	x	x
3.2.3.2.4.				x	x	x	x	x	x	x
3.2.3.2.5.				x	x	x	x	x	x	x
3.2.3.3.				x	x	x	x	x	x	x

Non-Kernel Security Related Software B5 Specification

3.2.4.1.	x x x x	x x x
3.2.4.2.	x x x x	x x x
3.2.4.3.	x x x x	x x x
3.2.4.4.	x x x x	x x x

Non-Kernel Security Related Software B5 Specification

5. PREPARATION FOR DELIVERY

This section is not applicable to this specification.

Non-Kernel Security Related Software B5 Specification

6. NOTES

This section is not applicable to this specification.

Non-Kernel Security Related Software B5 Specification

Appendix A: MAJOR CHANGES TO THE DESIGN

The main body of this document outlines the intended design of the NKSR portion of KSOS. During the implementation process, however, some aspects of the design have changed. This appendix gives a brief outline of the major changes in the design. Comments are given below by CPC and are in order of the CPC number.

The basic combined functionality for the Secure Initiator (CPC 3.1.1) and the Secure Server (CPC 3.1.2) has not changed between the previous version of the B5's and the current version. This functionality, however, was repropor-tioned between the two processes and the B5 sections were written to reflect this change.

At this time, Login (CPC 3.1.3) does not have the mechanism for locking out terminals.

The Password Modification Function (PMF), like Login and Logout, shall be contained within the code of the Secure Server. This function, however, is not currently implemented.

Level Preserving Copy (CPC 3.1.8) and Level Preserving Print (CPC 3.1.9) are implemented but not yet debugged or tested.

The status of the Network Daemon (CPC 3.2.1) can be found in the final report for KSOS.

Some changes were made to the checks Mount (CPC 3.2.5) performs. Mount now checks that the mount on entry has not been mounted on already. The other change is that when Mount checks (using SMXflow) that the user has flow, in the case of security, to the file system to be mounted. In the case of integrity, the flow check is made as if the user had the privilege to violate the integrity *-property.

System Shutdown (CPC 3.2.8) has not been implemented. Currently shutdown consists of pressing the halt switch.

Secure Mail (CPC 3.2.9) needs to be redesigned to take into account changes made to the Kernel design (in particular, those changes dealing with the set uid problem).

The B5's for the UNIX Directory Manager (CPC 3.2.11) are located in the emulator B5's for historical reasons.

Make File System (or File System Initialization) (CPC 3.3.3) is implemented as three programs: the Pack Initialization program (PKI), the Extent Initialization program (EXI), and the Boot Copy program (BTCP).

The Directory Consistency Check (CPC 3.3.4) and the Kernel Name to Pathname Mapping (CPC 3.3.6) have not been implemented.

Modify File System Control Entry (CPC 3.3.7) does not print warning messages indicating the seriousness of the proposed changes. It also does not

Non-Kernel Security Related Software B5 Specification

remove the file system from the immigration officer data base. If this function were to be implemented, it could cause problems when trying to make modifications to the root file system. The problem lies in the fact that the root file system must be mounted to run the immigration officer and to access the immigration officer data base.

The User Control Data Base Editor (CPC 3.4.1) is implemented as two programs: a program to change user information, and a program to change group information. Also, there is no `K_secure_terminal_lock_call` as is mentioned.

At this time, all messages sent to the Audit Capture Process (3.4.2) are kept. They are not marked as to whether they were sent via the signal mechanism or the inter-process communication (ipc) mechanism. In addition, the messages are not reformatted, though a time stamp is appended.

A program called the Audit Capture Operator Interface (ACP_OP) is used to send messages from the Operator to the Audit Capture Process. These messages cause the Audit Capture Process to report the name of the current audit file, to switch to a new audit file, or to place the audit messages out on the console instead of a file (or vice versa). In addition, the interface will remove or print any audit file indicated. When printing the audit file, the interface will format the messages in a readable form.

The Privilege Control Process (CPC 3.4.3) has not been implemented. Its functionality, however, is performed by the Modify File System Control Entry (CPC 3.3.7).

The list of possible privileges which could be given or withdrawn by the process has changed. Below is a brief description of each privilege currently permitted:

- a. The ability to violate the simple security property. (This privilege is not intended to be used. Originally, it was include merely for completeness. It was found, however, to be useful during implementation to circumvent certain types of bugs and allow development to continue without having to wait for the bug to be fixed.)
- b. The ability to violate the simple integrity property.
- c. The ability to violate the *-property for security.
- d. The ability to violate the *-property for integrity.
- e. The ability to violate the security compartments.
- f. The ability to violate discretionary access.
- g. The ability to use the `K_link` and `K_unlink` primitives.
- h. The ability to use the `K_mount` and `K_unmount` primitives.
- i. The ability to change the following pieces of type independent information about an object.

Non-Kernel Security Related Software B5 Specification

1. The security level of the object.
 2. The integrity level of the object.
 3. The security compartments of the object.
 4. The owner of the object.
 5. The group of the object.
-
- j. The ability to save the process image in the swapping area.
 - k. The ability to lock a segment in core.
 - l. The ability to use the K_signal primitive.
 - m. The ability to use the K_halt primitive.
 - n. The ability to do volume valids through the K_device_function primitive.
 - o. The ability to walk through the process table.
 - p. The ability to use the execute bit to determine if one has read access to the object (used, for example, when searching directory structures).
 - q. The ability to change the terminal path.
 - r. The ability to change characteristics of a terminal; e.g., the speed or the parity of a terminal.
 - s. The ability to grant or revoke privileges using the K_set_priv primitive.

The Immigration Officer Process (CPC 3.4.4) is not implemented in full. It was planned to take the File System Restore program (CPC 3.2.2) and add the necessary checks of the file system and the handling of the immigration officer data base to it.

```

1  $("      MODULE          System Operation Services Specs
2              (Version 1.8, 11:48:07)
3              Module Name (NSO.specs)
4              File Name (/kr/sunix/specs/nksr/s.NSO.specs)
5          CONTENTS
6          TYPE          SPECIAL.specifications
7          LAST CHANGED  4/21/81,11:48:07
8  ")
9
10 MODULE SOS
11     TYPES
12
13 nonDisType: STRUCT_OF(
14         INTEGER securityLevel; SET_OF securityCat securityCatS;
15         INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
16     $(integrityCat is typically the null set)
17 daType: SET_OF daMode;
18 modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
19 tiiStruct: STRUCT_OF(nonDisType nd;
20         modeStruct da; INTEGER owner, group; SET_OF privType priv);
21 DBMEntry: STRUCT_OF(VECTOR_OF CHAR fileSysName;
22         VECTOR_OF CHAR mountOnName; seid mountOnSeid; fileNsp slot;
23         BOOLEAN readOnly);
24     $("In addition to the above fields, the implementation also includes
25     a timestamp, omitted from these specs for simplicity,
26     and TIIinfo(fileSysName), which is omitted here since
27     there is no need to include TIIinfo() as a specific field
28     in a STRUCT within the specs.")
29
30 DDBentry: STRUCT_OF(nonDisType maxLevel;
31         BOOLEAN unloaded,
32         inUse);
33
34 deviceProfileEntry: STRUCT_OF(seid devSeid; nonDisType maxLevel);
35
36     $(from fca)
37 fileNsp: {129 .. 255};
38 globalData: STRUCT_OF(INTEGER linkCount, openCount, timeLastMod; seid subtype;
39         BOOLEAN openAtCrash);
40     $(state information for all openable objects)
41
42 mountTableEntry: STRUCT_OF(seid rootSeid; BOOLEAN readOnly;
43         tiiStruct devTii; globalData devG1);
44
45     PARAMETERS
46
47 INTEGER unassignedUserId $(For use by deassign);
48 INTEGER operatorIntegrityLevel;
49 INTEGER maxSizeDBME;
50 seid nullSeid;
51 fileNsp rootFileNsp; $(nullSeid and rootFileNsp are for initializing DBME)
52 SET_OF deviceProfileEntry deviceProfile;
53     $(deviceProfile is determined at system generation time;
54     for each "allowable" device devSeid, it provides the maximum
55     level of the device. Used in initializing DDB below.)
56 SET_OF VECTOR_OF CHAR DBIOset;

```

```

57      $(DBIOset is determined at system generation time; contains
58      the name of all extents that are allowed to be mounted. Used
59      here in initializing DBIO.)
60
61      DEFINITIONS
62
63      BOOLEAN LE(nonDisType ans,x) IS
64          ans.securityLevel <= x.securityLevel          AND
65          ans.securityCatS SUBSET x.securityCatS          AND
66          ans.integrityLevel <= x.integrityLevel          AND
67          ans.integrityCatS SUBSET x.integrityCatS;
68      $(The basic ordering relation between two nonDisTypes)
69
70      EXTERNALREFS
71
72      FROM smx:
73      INTEGER SENlowLevel;
74      INTEGER SENhighLevel;
75      seid: DESIGNATOR;
76      secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
77                          tExtent, tNull};
78      privType: {
79          privFileUpdateStatus,    privLink,          privLockSeg,
80          privModifyPriv,          privMount,
81          privSetFileLevel,        privSetSegProcLevel,
82          privStickySeg,           privTerminalLock,
83          privViolSimpSecurity,     privViolStarSecurity,
84          privViolSimpIntegrity,    privViolStarIntegrity,
85          privViolDiscrAccess,      privSignal,        privWalkPTable,
86          privHalt,                 privKernelCall,    privViolCompartments,
87          privRealizeExecPermissions;
88
89      daMode: {daRead, daWrite, daExecute};
90      securityCat: DESIGNATOR;
91      integrityCat: DESIGNATOR;
92      domainType: {userDomain, supervisorDomain};
93
94      VFUN TIIinfo(seid s) -> tiiStruct st;
95      VFUN SMXflow(seid pSeid, oSeid; daType da) -> BOOLEAN b;          $(SMXflow)
96      VFUN SMXdap(seid pSeid, oSeid; daType da) -> BOOLEAN b;
97      VFUN SENseidNsp(seid anySeid) -> INTEGER nsp;
98      VFUN SENseidToInt(seid anySeid) -> INTEGER i;
99      VFUN SENseidType(seid s) -> secureEntityType set;          $(SENseidType)
100     VFUN SMXhasPriv(seid pSeid; privType privl) -> BOOLEAN b;          $(SMXhasPriv)
101
102     FROM fca:
103     VFUN FCAinfo(seid fSeid) -> globalData gl;          $(FCAinfo)
104     VFUN FCAmountTable(fileNsp nsp) -> mountTableEntry mte; $(FCAmountTable)
105
106     FROM udm:
107     OFUN mount(seid pdirSeid,nspSeid; VECTOR OF CHAR name);
108     OFUN unmount(seid pdirSeid; VECTOR OF CHAR name);
109
110     FROM ker:
111     OVFUN K_mount(seid dev; BOOLEAN readOnly) -> fileNsp newNsp; $(K_mount)
112     OFUN K_unmount(seid devSeid);

```



```

113
114
115
116     FUNCTIONS
117
118 $(----- VFUNS -----)
119 VFUN DDB(seid dSeid) -> DDBentry ans;
120     $(Device Data Base)
121     HIDDEN;
122     INITIALLY
123     ans =
124     (LET deviceProfileEntry x |( x INSET deviceProfile
125     AND x.devSeid = dSeid)
126     IN
127     IF x ~= ?
128     THEN
129     STRUCT(x.maxLevel, FALSE, FALSE)
130     ELSE
131     ?);
132
133 VFUN DBIO() -> SET OF VECTOR OF CHAR ans;
134     $(Data Base Immigration Officer - The name of the extent you
135     are trying to mount in nksrMount must be in DBIO)
136     HIDDEN;
137     INITIALLY
138     ans = DBIOset;
139
140 VFUN DBME() -> SET OF DBMEentry ans;
141     $(Data base of mounted extents.
142     Initialized in the implementation by Secure Initiator; the specs
143     below simulate this initialization. Entries are added
144     by nksrMount, deleted by nksrUnmount)
145
146     HIDDEN;
147     INITIALLY
148     ans = {STRUCT(
149     "/",      $(root fileSysName)
150     "",      $(null mountOnName)
151     nullSeid, $(null mountOnSeid)
152     rootFileNsp,
153     FALSE
154     )};
155
156
157 $(----- OFUNS -----)
158
159
160 OFUN assign(seid devSeid)[seid pSeid];
161
162     DEFINITIONS
163     tiiStruct ptii IS TIIinfo(pSeid);
164
165     EXCEPTIONS
166     notDevice: SENseidType(devSeid) ~= tDevice;
167     badDev: DDB(devSeid) = ?;
168     badLevel: NOT LE(ptii.nd,DDB(devSeid).maxLevel);

```

```

169  inUse: DDB(devSeid).inUse;
170
171  EFFECTS
172  `TIIinfo(devSeid) = STRUCT(
173      ptii.nd,
174      STRUCT({daRead,daWrite},{},{}), $(read/write by owner only)
175      SENseidToInt(pSeid), $(new owner is pSeid)
176      ptii.group,
177      ptii.priv);
178  `DDB(devSeid) = STRUCT(
179      DDB(devSeid).maxLevel,
180      TRUE, $(device is now "unloaded")
181      TRUE); $(device is now in use)
182
183  OFUN deassign(seid devSeid) [seid pSeid];
184
185  DEFINITIONS
186  tiiStruct dtii IS TIIinfo(devSeid);
187  tiiStruct ptii IS TIIinfo(pSeid);
188
189  EXCEPTIONS
190  notDevice: SENseidType(devSeid) ~= tDevice;
191  badDev: DDB(devSeid) = ?;
192  badOwner: dtii.owner ~= SENseidToInt(pSeid) AND
193      ptii.nd.integrityLevel < operatorIntegrityLevel;
194      $(Either the invoking process owns the device, or is
195      at or above operator integrity level)
196
197  EFFECTS
198  `TIIinfo(devSeid) = STRUCT(
199      STRUCT(SENhighLevel,{},SENhighLevel,{}),
200      STRUCT({},{}),{}),
201      unassignedUserId,
202      unassignedUserId,
203      {}
204      );
205
206  `DDB(devSeid) = STRUCT(
207      DDB(devSeid).maxLevel,
208      DDB(devSeid).unloaded,
209      FALSE
210      );
211
212  OFUN nksrMount( VECTOR_OF CHAR mountFromName;
213      seid mountFromSeid;
214      VECTOR_OF CHAR mountOnName;
215      seid mountOnSeid;
216      BOOLEAN readOnly
217      ) [seid pSeid];
218
219  $(The parameters mountFromSeid and mountOnSeid are given explicitly
220  in these specs for convenience. In the implementation, these seids
221  are determined from mountFromName, drive number, and mountOnName resp.)
222
223  $("Between the kernel and the nksr, there are three mount functions.
224  K-mount performs physical mounting; mount in the directory manager

```

```

225 performs a simple mapping in the directory VFUN; the present
226 nksrMount utilizes the functionality of the other two mounts.
227 As a result of nksrMount, an extent is now interpreted as a file system
228 by the directory manager.")
229
230 DEFINITIONS
231 fileNsp openSlot IS
232     SOME fileNsp freeMountNsp | FCAMountTable(freeMountNsp)=?;
233 $(openSlot simulates the value returned by K_mount;
234 the related value nspSeid, defined next, is used in the
235 "call" to mount in the EFFECTS section below.)
236
237 seid nspSeid IS
238     SOME seid s | SENseidNsp(s) = openSlot;
239     $(Need a seid, rather than simply an nsp component to pass to mount)
240
241 EXCEPTIONS
242 alreadyMounted: EXISTS DBMEentry x INSET DBME() :
243     x.fileSysName = mountFromName;
244 badExtent: NOT mountFromName INSET DBIO(): $(The extent you are trying
245     to mount must be in the Immigration Officer database, DBIO)
246 DBMEfull: CARDINALITY(DBME()) = maxSizeDBME;
247 badRWAccess: NOT readOnly AND NOT SMXflow(pSeid, mountFromSeid,
248     {daRead,daWrite});
249 badRAccess: readOnly AND NOT SMXflow(pSeid,mountFromSeid,{daRead});
250 badWAccess: NOT SMXflow(pSeid, mountOnSeid, {daWrite});
251 badRWdap: NOT readOnly AND NOT SMXdap(pSeid, mountFromSeid,
252     {daRead,daWrite});
253 badRdap: readOnly AND NOT SMXdap(pSeid,mountFromSeid,{daRead});
254 badWdap: NOT SMXdap(pSeid, mountOnSeid, {daWrite});
255 notStarSec: NOT SMXhasPriv(pSeid,privViolStarSecurity);
256 $(Exceptions from FCAMount follow)
257 XbadPriv: NOT SMXhasPriv(pSeid,privMount);
258 XnoFile: FCAinfo(mountFromSeid) = ?; $(Successfully passing the
259     badExtent check above implies that this check will
260     be passed successfully)
261 XnoClass: SENseidType(mountFromSeid) ~= tExtent;$(same comment as XnoFile)
262 XmntSpace: RESOURCE_ERROR;
263
264 EXCEPTIONS_OF mount(mountOnSeid,nspSeid,mountOnName):
265
266 EFFECTS
267 openSlot = EFFECTS_OF K_mount(mountFromSeid,readOnly);
268 EFFECTS_OF mount(mountOnSeid,nspSeid,mountOnName);
269 ^DBME() = DBME() UNION {STRUCT(
270     mountFromName,
271     mountOnName,
272     mountOnSeid,
273     openSlot,
274     readOnly
275     )};
276
277
278 OFUN nksrUnmount(seid fSeid) [seid pSeid];
279     $(fSeid corresponds to mountOnSeid from nksrMount)
280

```

```

281      DEFINITIONS
282      DBMEentry x IS SOME DBMEentry y | (y INSET DBME() AND
283        y.mountOnSeid = fSeid);
284      $(It is assumed that mountOnSeid, as a field within DBMEentry,
285        is a key, i.e., uniquely determines the entire entry y.)
286      BOOLEAN readOnly IS x.readOnly;
287      fileNsp fsNsp IS x.slot;
288
289      EXCEPTIONS
290
291      NotalreadyMounted: x = ?;
292      badWAccess: NOT SMXflow(pSeid, fSeid, {daWrite});
293      badWdap: NOT SMXdap(pSeid, fSeid, {daWrite});
294      notStarSec: NOT SMXhasPriv(pSeid, privViolStarSecurity);
295      XbadPriv: NOT SMXhasPriv(pSeid, privMount);
296      XnoXXX: FCAmountTable(fsNsp) = ?;
297      XnoTranquil:
298        EXISTS seid fSeid | SENseidNsp(fSeid) = fsNsp
299          : FCAinfo(fSeid).openCount > 0;
300      EXCEPTIONS_OF unmount(fSeid, x.mountOnName);
301
302      EFFECTS
303
304      EFFECTS_OF K_unmount(fSeid);
305      EFFECTS_OF unmount(fSeid, x.mountOnName);
306      ^DBME() = DBME() DIFF {x};
307
308
309      END_MODULE

```

```

1 $(" MODULE          Secure User Services Specs
2                               Version (1.7,14:52:13)
3                               Module Name (NSU.specs)
4                               File Name (/kr/sunx/specs/nksr/s.NSU.specs)
5 CONTENTS                  Secure Server, Login, Logout, ChangeAccess Level,
6                               ChangeGroup, FileAccessModifier
7 TYPE                      SPECIAL.specifications
8 LAST CHANGED              11/21/80,14:52:13
9 ")
10
11
12 MODULE SUS
13
14 TYPES
15
16 nonDisType: STRUCT_OF(
17     INTEGER securityLevel; SET_OF securityCat securityCatS;
18     INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
19 $(integrityCat is typically the null set)
20 daType: SET_OF daMode;
21 modeStruct: STRUCT_OF(daType ownerMode, groupMode, allMode);
22 tiStruct: STRUCT_OF(nonDisType nd;
23     modeStruct da; INTEGER owner, group; SET_OF privType priv);
24
25
26 PARAMETERS
27
28 INTEGER sysHighLevel $(Highest possible security and integrity level);
29 SET_OF securityCat allSecCats $(Full universe of securityCatS) ;
30 SET_OF integrityCat allIntCats $(Full universe of integrityCatS) ;
31
32
33
34 DEFINITIONS
35
36 nonDisType minimum(SET_OF nonDisType sd) IS
37     SOME nonDisType ans | ans INSET sd                                AND
38     (FORALL nonDisType x INSET sd:
39         ans.securityLevel <= x.securityLevel                        AND
40         (FORALL securityCat a INSET ans.securityCatS : a INSET
41             x.securityCatS)                                          AND
42         ans.integrityLevel <= x.integrityLevel);
43
44 BOOLEAN LE(nonDisType ans,x) IS
45     ans.securityLevel <= x.securityLevel                            AND
46     ans.securityCatS SUBSET x.securityCatS                          AND
47     ans.integrityLevel <= x.integrityLevel                          AND
48     ans.integrityCatS SUBSET x.integrityCatS;
49 $(The basic ordering relation between two nonDisTypes)
50
51
52
53 EXTERNALREFS
54
55 FROM smx:
56 seId: DESIGNATOR;

```

```

57 secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
58                     tExtent, tNull};
59 privType: {
60     privFileUpdateStatus,    privLink,        privLock;
61     privModifyPriv,          privMount,
62     privSetFileLevel,        privSetSegProcLevel,
63     privStickySeg,           privTerminalLock,
64     privViolSimpSecurity,     privViolStarSecurity,
65     privViolSimpIntegrity,    privViolStarIntegrity,
66     privViolDiscrAccess,      privSignal,      privWalkPTable,
67     privHalt,                 privKernelCall, privViolCompartments,
68     privRealizeExecPermissions};
69
70 daMode: {daRead, daWrite, daExecute};
71 securityCat: DESIGNATOR;
72 integrityCat: DESIGNATOR;
73 domainType: {userDomain, supervisorDomain};
74
75 VFUN TIIInfo(seid s) -> tiiStruct st;
76 VFUN SMXflow(seid pSeid, oSeid; daType da) -> BOOLEAN b:          $(SMXflow)
77 VFUN FCATerminalPathSet(seid terminalGroup) -> SET_OF seid ss;
78 VFUN FCAcurrentPath(seid terminalGroup) -> seid s;
79     $(Note that the parameter types of the two FCA VFUNS has been changed
80     from a separate DESIGNATOR type to a seid.
81     These changes must also be made in the kernel specs.)
82 VFUN SENseidNsp(seid anySeid) -> INTEGER nsp;
83
84
85     ASSERTIONS
86
87 TRUE;    $(Placeholder for the ASSERTIONS paragraph.)
88
89
90     FUNCTIONS
91
92 $(-----VFUNS-----)
93
94 $(The important VFUNS are:
95  UAADB (User Authentication Access Data Base)
96  TPDB (Terminal Profile Data Base)
97  SDB (System Data Base)
98  GADB (Group Authorization Data Base)
99  FCATerminalPathSet
100 FCAcurrentPath
101 terminalUser (for a given terminal, tells the user, as given by
102     the most recent Login)
103 ptbl (a table showing, for a given process seid, the terminal
104     and terminal path in use by that seid)
105 auditTrail (for sending audit messages)
106 securePath (path used when ATTN key is struck on a terminal)
107
108
109 ")
110
111 VFUN UAADB(VECTOR OF CHAR userName) -> STRUCT_OF(
112     VECTOR_OF CHAR passWord;

```

```

113     INTEGER consecutiveUnsuccessfulLogins;
114     BOOLEAN loginsPermitted;
115     seid initialProcess;
116     tiiStruct initialLevel;
117     nonDisType maxLevel) ans;
118     $("Life cycle of UAADB: logIn reads the following:
119     passWord, loginsPermitted, initialProcess, initialLevel, maxLevel ;
120     logIn updates consecutiveUnsuccessfulLogins")
121     HIDDEN;
122     INITIALLY
123     TRUE;
124     $(***Perhaps there are some initial restrictions?)
125
126 VFUN TPDB(seid terminalGroup ) -> STRUCT_OF(
127     BOOLEAN locked,
128         configured,
129         loggedIn;
130     nonDisType maxLevel) ans :
131
132     $(Life cycle of TPDB: logIn reads the following fields:
133     locked, configured, loggedIn, maxLevel.
134     changeAccessLevel reads maxLevel; logIn updates loggedIn.)
135
136     HIDDEN;
137     INITIALLY
138     TRUE;
139
140 VFUN SDB() -> STRUCT_OF(
141     nonDisType maxLevel) x; $(given as STRUCT to allow for more fields later)
142     $(read by changeAccessLevel, logIn, fileAccessModifier)
143     HIDDEN;
144     INITIALLY
145     TRUE;
146
147 VFUN GADB(INTEGER x) -> STRUCT_OF(
148     nonDisType maxLevel;
149     VECTOR_OF CHAR password) y ;
150     $(read by changeGroup, changeAccessLevel, fileAccessModifier)
151     HIDDEN;
152     INITIALLY
153     TRUE;
154
155 VFUN terminalUser(seid terminalGroup) -> VECTOR_OF CHAR user;
156     $(initialized by logIn; finalized by logOut; read by fileAccessModifier)
157     HIDDEN;
158     INITIALLY
159     FORALL seid x : terminalUser(x) =?;
160
161
162 VFUN ptbl(seid pSeid) -> STRUCT_OF(
163     seid t,    $("terminal")
164     p) x ;    $("path used by pSeid on terminal t")
165     HIDDEN;
166     INITIALLY
167     FORALL seid x: ptbl(x) =? ;
168     $("ptbl is initialized by the kernel. For simplicity in these

```

```

169 specs, however, we present only the above initial'ation.
170 ptbl is updated by logIn; read by logOut; and finalized by
171 logOut; In the implementation there is the additional complexity
172 that the logIn process might create new processes, which are
173 also added to ptbl. We ignore this complexity in these specs")
174
175 VFUN auditTrail() -> VECTOR_OF VECTOR_OF CHAR x;
176     HIDDEN;
177     INITIALLY
178         LENGTH(x) = 0;
179     $("auditTrail is a sequence of messages sent to the audit capture process.
180     In the implementation, each message is fixed-length and of a given
181     format. Such restrictions are ignored here; the main purpose of
182     including auditTrail in these specs is to show when such messages
183     are sent")
184
185 VFUN securePath(seid terminalGroup) -> seid x;
186     HIDDEN;
187     INITIALLY
188         x INSET FCATerminalPathSet(terminalGroup);
189
190 VFUN charsToInt(VECTOR_OF CHAR x) -> INTEGER y;
191     HIDDEN;
192     INITIALLY
193         TRUE;
194     $(A minor conversion function used by fileAccessModifier)
195
196
197
198 $(-----OFUNS-----)
199
200 $( "SecureInitiator and SecureServer are not given explicitly
201 in these specs. Instead, the main events they handle,
202 namely, detecting that the attention key is struck, detecting
203 that a hangup has occurred destroying the physical connection between
204 a terminal and the cpu, and operator requested logout -- are given
205 by the OFUNS ATTN, HANGUP, and OPREQLOGOUT. The remainder of the
206 functionality of the SecureServer is that of a command interpreter.
207 In line with the HDM model of computation, the semantics for the
208 individual commands is given below, but the notion of an explicit
209 command interpreter is not given.")
210
211 OFUN ATTN()[seid terminalGroup];
212     $(In response to the attention key being struck on the terminal)
213
214     EFFECTS
215
216     ^FCACurrentPath(terminalGroup) = securePath(terminalGroup);
217     $("Since TIIinfo(securePath(t)) never changes, there is no need
218     to set it here.")
219
220
221 OFUN HANGUP()[seid terminalGroup];
222     $(In response to the detection of a broken terminal connection)
223
224     DEFINITIONS

```



```

225   seid t IS terminalGroup;
226
227   EXCEPTIONS
228   notLoggedIn: NOT TPDB(t).loggedIn;
229
230   EFFECTS
231   FORALL seid x | ptbl(x).t = t AND
232     ptbl(x).p ~= securePath(t) :
233     'TIIinfo(x) = ? AND 'ptbl(x) = ?;
234     $(Wipe out all processes which are using a non-secure path on the
235     effected terminal.)
236
237   'terminalUser(t) = ?;
238   'TPDB(t).loggedIn = FALSE;
239   FORALL seid x INSET FCATerminalPathSet(t):
240     x ~= securePath(t) =>
241     'TIIinfo(x) = ?;
242   'FCACurrentPath(t) = ?;
243   $(Wipe out all paths, except secure path, on the effected terminal.)
244
245   OFUN OPREQLOGOUT(seid terminalGroup)[seid pSeid]:
246     $(Note that terminalGroup is given explicitly for an operator
247     requested logout, in contrast to normal logout or hangup)
248
249   DEFINITIONS
250   seid t IS terminalGroup;
251
252   EXCEPTIONS
253   notLoggedIn: NOT TPDB(t).loggedIn;
254   badpSeid: TIIinfo(pSeid).nd ~= STRUCT(sysHighLevel, allSecCats, 'suighLevel,
255     allIntCats);
256   $(badpSeid indicates that somebody other than the operator is attempting
257   to invoke OPREQLOGOUT)
258
259   EFFECTS
260   FORALL seid x | ptbl(x).t = t AND
261     ptbl(x).p ~= securePath(t) :
262     'TIIinfo(x) = ? AND 'ptbl(x) = ?;
263
264   'terminalUser(t) = ?;
265   'TPDB(t).loggedIn = FALSE;
266   FORALL seid x INSET FCATerminalPathSet(t):
267     x ~= securePath(t) =>
268     'TIIinfo(x) = ?;
269   'FCACurrentPath(t) = ?;
270   'auditTrail() = VECTOR (FOR i FROM 1 TO LENGTH(auditTrail())+1:
271     IF i <= LENGTH(auditTrail()) THEN auditTrail()[i]
272     ELSE "operator requested logout" );
273
274
275   OFUN logIn(VECTOR_OF CHAR userName,passWord) [seid terminalGroup];
276
277   DEFINITIONS
278   seid t IS terminalGroup;
279   seid p IS SOME seid q | q INSET FCATerminalPathSet(t) AND
280     TIIinfo(p)=?;

```

```

281    $("Note that p cannot be securePath(t)")
282    ttiStruct s IS UAADB(userName).initialLevel;
283    nonDisType ndLevel IS minimum({
284        s.nd,
285        TPDB(t).maxLevel,
286        UAADB(userName).maxLevel,
287        GADB(UAADB(userName).initialLevel.group).maxLevel,
288        SDB().maxLevel});
289    ttiStruct startLevel IS STRUCT(ndLevel, s.da, s.owner, s.group, s.priv);
290
291    EXCEPTIONS
292    badt: TPDB(t) = ?;
293    badTPDBdata: TPDB(t).locked OR NOT TPDB(t).configured
294        OR TPDB(t).loggedIn OR NOT UAADB(userName).loginsPermitted;
295    undefUser: UAADB(userName) = ?;
296    noPath: p = ?;
297    badLevel: ndLevel = ?;          $(This should never happen)
298
299    EFFECTS
300    IF (passWord ~= UAADB(userName).passWord) THEN
301        UAADB(userName) = STRUCT(UAADB(userName).passWord,
302            UAADB(userName).consecutiveUnsuccessfulLogins +1,
303            UAADB(userName).loginsPermitted,
304            UAADB(userName).initialProcess,
305            UAADB(userName).initialLevel,
306            UAADB(userName).maxLevel)
307    ELSE
308        UAADB(userName) = STRUCT(UAADB(userName).passWord,
309            0,
310            UAADB(userName).loginsPermitted,
311            UAADB(userName).initialProcess,
312            UAADB(userName).initialLevel,
313            UAADB(userName).maxLevel)
314        TTIinfo(p) = startLevel
315        TTIinfo(UAADB(userName).initialProcess) = startLevel
316        terminalUser(t) = userName
317        TPDB(t).loggedIn = TRUE
318        FCACurrentPath(t) = p
319        ptbl(UAADB(userName).initialProcess) = STRUCT(t,p) ;
320        auditTrail() = VECTOR (FOR i FROM 1 TO LENGTH(auditTrail())+1:
321            IF i <= LENGTH(auditTrail()) THEN auditTrail()[i]
322            ELSE "login attempted" );
323
324
325
326    OFUN logOut()[seid terminalGroup];
327
328    DEFINITIONS
329    seid t IS terminalGroup;
330
331    EXCEPTIONS
332    notLoggedIn: NOT TPDB(t).loggedIn;
333
334    EFFECTS
335    FORALL seid x | ptbl(x).t = t AND
336        ptbl(x).p ~= securePath(t) ;

```

```

337      ^TIIinfo(x) = ? AND ^ptbl(x) = ?;
338      ^terminalUser(t) = ?;
339      ^TPDB(t).loggedIn = FALSE;
340      FORALL seid x INSET FCATerminalPathSet(t):
341          x ^= securePath(t) =>
342              ^TIIinfo(x) = ?;
343
344      ^FCACurrentPath(t) = ?;
345      ^auditTrail() = VECTOR (FOR i FROM 1 TO LENGTH(auditTrail())+1:
346          IF i <= LENGTH(auditTrail()) THEN auditTrail()[i]
347          ELSE "logout" );
348
349
350      OFUN changeAccessLevel(seid terminalGroup; tiiStruct desiredLevel) ;
351
352      DEFINITIONS
353      seid t IS terminalGroup;
354      seid q IS SOME seid p |
355          p INSET FCATerminalPathSet(t) DIFF {securePath(t)} AND
356          TIIinfo(p).nd = desiredLevel.nd AND
357          TIIinfo(p).owner = desiredLevel.owner AND
358          TIIinfo(p).group = desiredLevel.group ;
359      seid r IS SOME seid p |
360          p INSET FCATerminalPathSet(t) DIFF {securePath(t)} AND
361          TIIinfo(p) = ?;
362      ronDisType ndLevel IS minimum({
363          TPDB(t).maxLevel,
364          SDB().maxLevel,
365          UAADB(terminalUser(t)).maxLevel,
366          GADB(TIIinfo(FCACurrentPath(t)).group).maxLevel
367      });
368      tiiStruct cl $(current level) IS
369          TIIinfo(FCACurrentPath(t));
370      tiiStruct startLevel IS STRUCT(ndLevel,cl.da,cl.owner,cl.group,cl.priv);
371
372      EXCEPTIONS
373      noPath: q = ? AND r = ?;
374      badLevel: NOT LE(desiredLevel.nd, ndLevel);
375
376      EFFECTS
377      IF q ^= ? THEN
378          ^FCACurrentPath(t) = q
379      ELSE
380          ^TIIinfo(r) = startLevel AND
381          ^TIIinfo(UAADB(terminalUser(t)).initialProcess) = startLevel AND
382          ^FCACurrentPath(t) = r ;
383      ^auditTrail() = VECTOR (FOR i FROM 1 TO LENGTH(auditTrail())+1:
384          IF i <= LENGTH(auditTrail()) THEN auditTrail()[i]
385          ELSE "change access level" );
386
387      OFUN changeGroup(seid terminalGroup; seid desiredGroup;
388          VECTOR OF CHAR pwdDesiredGroup) [seid pSeid] ;
389      $(pSeid is the seid in whose behalf the change is being made)
390
391      DEFINITIONS
392      seid t IS terminalGroup;

```

```

393  tiiStruct ptii IS TIIinfo(pSeid);
394  seid q IS SOME seid p |
395    p INSET FCATerminalPathSet(t) DIFF {securePath(t)} AND
396    TIIinfo(p).nd = TIIinfo(desiredGroup).nd
397    TIIinfo(p).owner = TIIinfo(desiredGroup).owner
398    TIIinfo(p).group = TIIinfo(desiredGroup).group ;
399  seid r IS SOME seid p |
400    p INSET FCATerminalPathSet(t) DIFF {securePath(t)} AND
401    TIIinfo(p) = ?;
402  nonDisType ndLevel IS minimum({
403    TPDB(t).maxLevel,
404    SDB().maxLevel,
405    UAADB(terminalUser(t)).maxLevel,
406    TIIinfo(FCACurrentPath(t)).nd
407  });
408  tiiStruct cl S(current level) IS
409    TIIinfo(FCACurrentPath(t));
410  tiiStruct startLevel IS STRUCT(cl.nd,cl.da,cl.owner,
411    SENseidNsp(desiredGroup),
412    cl.priv);
413
414  EXCEPTIONS
415  badt: TPDB(t) = ?;
416  noPath: q = ? AND r = ?;
417  badLevel: NOT LE(GADB(SENseidNsp(desiredGroup)).maxLevel,ndLevel);
418  badGroup: GADB(SENseidNsp(desiredGroup)) = ?;
419  badPassWord: pwdDesiredGroup ~=
420    GADB(SENseidNsp(desiredGroup)).password;
421
422  EFFECTS
423  ^TIIinfo(pSeid) = startLevel;
424  IF q ~= ? THEN
425    ^FCACurrentPath(t) = q
426  ELSE
427    ^TIIinfo(r) = startLevel
428    ^TIIinfo(UAADB(terminalUser(t)).initialProcess) = startLevel
429    ^FCACurrentPath(t) = r ;
430  ^auditTrail() = VECTOR (FOR i FROM 1 TO LENGTH(auditTrail())+1:
431    IF i <= LENGTH(auditTrail()) THEN auditTrail()[i]
432    ELSE "change group" );
433
434
435  OFUN fileAccessModifier(seid terminalGroup, fSeid; nonDisType desirednd:
436    modeStruct da);
437
438  DEFINITIONS
439  seid t IS terminalGroup;
440  tiiStruct ftii IS TIIinfo(fSeid);
441
442  EXCEPTIONS
443  badt: TPDB(t) = ?;
444  badFile: TIIinfo(fSeid) = ?;
445  badOwner: ftii.owner ~= charsToInt(terminalUser(t));
446    $(User must own file he is trving to modify)
447  badFileLevel: ftii.nd ~=
448    minimum

```

```

449     ftii.nd,
450     TPDB(t).maxLevel,
451     SDB().maxLevel,
452     UAADB(terminalUser(t)).maxLevel,
453     GADB(UAADB(terminalUser(t)).initialLevel.group).maxLevel
454   });
455   badDesiredLevel: desirednd ~=
456     minimum({
457       desirednd,
458       TPDB(t).maxLevel,
459       SDB().maxLevel,
460       UAADB(terminalUser(t)).maxLevel,
461       GADB(UAADB(terminalUser(t)).initialLevel.group).maxLevel
462     });
463
464
465   EFFECTS
466   ^TIIinfo(fSeid) = STRUCT(desirednd, da, ftii.owner, ftii.group,
467     ftii.priv);
468   ^auditTrail() = VECTOR (FOR i FROM 1 TO LENGTH(auditTrail())+1:
469     IF i <= LENGTH(auditTrail()) THEN auditTrail()[i]
470     ELSE "file fSeid modified to desired nd" );
471
472
473 END_MODULE

```

```
1 MODULE alc          $(access level control)
2 $("
3 COPYRIGHT: 1978, by Ford Aerospace and Communications Corp.
4 ADDRESS:   Ford Aerospace and Communications Corp.
5           Western Development Laboratories
6           3939 Fabian Way
7           Palo Alto, California 94303
8           Attention: Software Technology Dept.
9           Mail Stop: V02
10
11 CPCI:  NKSr
12 CPC:
13
14 MODULE NAME:  alc  access level control
15 FUNCTION:
16 ASSUMPTIONS:
17 HISTORY:
18   AUTHOR:   Tom Berson
19   VERSION:  1.1 of 10/18/78
20   MODULE TYPE: SPECIAL Specifications
21 SPECIAL NOTES:
22 ")
23
24
25 TYPES
26
27 accessLevel: STRUCT_OF(INTEGER securityLevel,integrityLevel;
28                        SET_OF securityCat sCatSet);
29
30 PARAMETERS
31
32 accessLevel ALCgroundLevel; $(level for terminals not logged in)
33 accessLevel ALCloginLevel;  $(initial level for all logins)
34
35
36 EXTERNALREFS
37
38   FROM smx:
39 securityCat: DESIGNATOR;
40 integrityCat: DESIGNATOR;
41
42
43 FUNCTIONS
44
45 VFUN ALCsetPermissible(accessLevel p,max)->BOOLEAN B;
46
47   DERIVATION
48   (p.securityLevel<=max.securityLevel)
49 AND (p.sCatSet SUBSET max.sCatSet)
50 AND (p.integrityLevel<=max.integrityLevel);
51
52
53 END_MODULE
```

```

1  $( " MODULE:      despooler.specs (Version 1.9, 14:13:09)
2                      Module Name (dsp.specs)
3                      File Name (/kr/sunix/specs/nksr/s.dsp.specs)
4      CONTENTS      despooler operations
5      TYPE:         SPECIAL specifications
6      LAST CHANGED:  11/21/80,14:13:09
7
8  " )
9
10 MODULE despooler
11
12 $( "The main purpose of the despooler is to periodically check
13 if files are spooled for printing, and if so, queue
14 them to be printed one at a time. One main interface is thus
15 to the spooler mechanism, which indicates the currently spooled
16 files in the VFUN spooledFiles. The spooler mechanism puts files
17 from the "outside world" into spooledFiles via the OFUN spool.
18
19 The other main interface of the despooler is to the "operator",
20 i.e., someone who sends the following sorts of commands to the
21 despooler: start, stop, die, setLevel. Start, stop, and die
22 force the despooler to look like a state transition machine, in that
23 printing of files can only occur from state started. The command
24 setLevel does not affect the state transition aspects, but merely
25 changes the set of files that will actually be printed.
26 Also, the command start can optionally specify a new level at which
27 the despooler will be started. In these specs, the new level will
28 always be given as a parameter of start. The level
29 of each file to be printed must be <= maxPrintLevel, which is a VFUN
30 modified as a result of interpreting the command setLevel or start. For each
31 of the four commands, the operator also sends an indication of whether
32 the command should be interpreted "immediately" or after the currently
33 queued files are all printed.
34
35 In addition to the "applications-oriented" tasks of queueing and
36 printing files, the interrupt-handling aspects of the operator commands
37 must also be indicated in these specs.
38
39 As a way of structuring these specs, the following six categories
40 of objects are useful:
41
42 I. Operator Commands - setLevel, start, stop, die.
43 II. Interrupt Buffers - Boolean VFUNS which are initially FALSE;
44     they are set to TRUE by OFUN issueCommand, and reset to FALSE
45     when the interrupt is actually processed by OFUN processInterrupts.
46     The Boolean VFUNS are: levbuf, startbuf, stopbuf, and diebuf,
47     corresponding respectively to the four operator commands above.
48     In addition, there is the buffer newLevbuf, indicating the new level
49     for a setLevel or start command.
50 III. ILPL events. Below are given abstract programs in ILPL for
51     the spooler, despooler, and operator. One of the structuring
52     mechanisms of ILPL is based on Zahn's "event-indicator" construct
53     (see Knuth's "Str. Pgmng with Gotos" article).
54     The ILPL events used in the programs below are:
55     startE, stopE, dieE, immInt ("immediate interrupt" ), Qempty.
56 IV. BOOLEAN VFUNS corresponding to the ILPL events.

```

```

57      VstartE, VstopE, VdieE. Initially FALSE; tested by the ILPL
58      programs to determine if the corresponding events should be
59      signalled. Set to TRUE by OFUN processInterrupts. Reset
60      to FALSE by OFUNS given in the next category, which are invoked
61      by the ILPL program prior to signalling the event.
62      (No specific VFUNS are needed for Qempty, which gets signalled
63      when LENGTH(queue) = 0; or for immInt, which gets signalled
64      when nowbuf is TRUE.
65      V. OFUNS to reset the BOOLEAN VFUNS in category IV to FALSE:
66          OstartE, OstopE, OdieE.
67          These OFUNS are invoked by the ILPL program, and reset the
68          corresponding VFUNS to FALSE.
69      VI. OFUNS and VFUNS appropriate to spooler/despooler functionality.
70          The following are the major ones:
71          VFUNS: spooledFiles - spooler puts files here; despooler
72                  takes files from here and puts them on queue.
73          queue - files wait here until they are printed.
74                  If a setlevel interrupt occurs, the queue may
75                  be cleared prior to printing the files. The
76                  files are still in spooledFiles.
77          output - simulates line printer output. New printed
78                  pages are added to the right of this vector.
79          maxPrintLevel - changed by setlevel. Only files <=
80                  to this level may be moved from the outside
81                  world to spooledFiles, or from spooledFiles
82                  to queue.
83          currentState - handles state transitions; used mainly
84                  by OFUN processInterrupts.
85
86          OFUNS:
87              issueCommand - receive command from operator, and set
88                  interrupt buffers appropriately.
89              processInterrupts - check the interrupt buffers
90                  and take appropriate action. Handle state
91                  transitions.
92              processQueue - dequeue and print the file at the
93                  beginning of the queue.
94              fillQueue - move files from spooledFiles to queue.
95              spool - move files from the outside world to spooledFiles.
96
97      -----
98
99      The ILPL program for the despooler is given next.
100
101      (cycle until a start command is received. This also sets maxPrintLevel.)
102      UNTIL dieE DO
103          UNTIL stopE,dieE DO
104              processInterrupts;
105              IF VstopE THEN
106                  OstopE;
107                  SIGNAL(stopE);
108              END IF
109              IF VdieE THEN
110                  OdieE;
111                  SIGNAL(dieE);
112              END IF


```



```

113     fillQueue;
114     UNTIL Qempty, immInt DO
115       IF LENGTH(queue) = 0 THEN;
116         SIGNAL(Qempty)
117       END IF
118       IF nowbuf() THEN;
119         SIGNAL(immInt);
120       END IF
121       processQelement;
122     THEN
123       ON Qempty: ;
124       ON immInt:
125         processInterrupts;
126         IF VstopE THEN
127           OstopE;
128           SIGNAL(stopE);
129         END IF
130         IF VdieE THEN
131           OdieE;
132           SIGNAL(dieE);
133         END IF
134       END
135     THEN
136       ON dieE: SIGNAL(dieE);
137       ON stopE:
138         UNTIL startE DO
139           processInterrupts;
140           IF VstartE() THEN
141             OstartE;
142             SIGNAL(startE);
143           END IF;
144         THEN
145           ON startE: ;
146         END;
147       END;
148     THEN
149     ON dieE: "abort";
150     END
151     -----
152     The abstract program for the spooler is:
153     UNTIL doomsday DO
154       spool(cFile,tFile);
155     END;
156     -----
157     The abstract program for the operator is:
158     UNTIL doomsday DO
159       issueCommand(cmd);
160     END;
161     -----
162     The operator is "supposed" to issue commands in accordance with
163     the following state transition diagram.
164     -----
165     setlevel (level)

```



183 processInterrupts makes sure that this diagram is obeyed.

```

184 -----
185
186 ")
187
188
189     TYPES
190 $(from smx)
191 nonDisType: STRUCT_OF(
192     INTEGER securityLevel; SET_OF securityCat securityCatS;
193     INTEGER integrityLevel; SET_OF integrityCat integrityCatS);
194 $(integrityCat is typically the null set)
195 daType: SET_OF daMode;
196 modeStruct: STRUCT_OF(daType ownerMode, groupMode*, allMode);
197 tiiStruct: STRUCT_OF(nonDisType nd;
198     modeStruct da; INTEGER owner, group; SET_OF privType priv);
199 $(from fca)
200 globalData: STRUCT_OF(INTEGER linkCount, openCount, timeLastMod; seid subtype;
201     BOOLEAN openAtCrash);
202 $(state information for all openable objects)
203
204
205 $(Types for despooler are given next.)
206 line: {VECTOR_OF CHAR x | LENGTH(x) <= maxLineLength};
207 pageImage: {VECTOR_OF line x | LENGTH(x) <= maxPageLength};
208 outStream: VECTOR_OF pageImage;
209 textImage: VECTOR_OF line;
210 $(variables of type textImage will be text files containing
211     an unbounded but finite sequence of lines.)
212 commands: {setLevel, stop, die, start};
213 $(The commands are IPC messages that can be received by
214     the despooler.)
215 dspStates: {started, stopped, dead};
216 marMsgType : {VECTOR_OF line x | LENGTH(x) = marginHeadingLength};
217 $(type of the security message that appears at top and bottom of
218     each output page.)
219
220     PARAMETERS
221
222     INTEGER maxPageLength $(Number of physical lines on a physical page) ;
223     INTEGER maxLineLength;
224     INTEGER marginHeadingLength; $(Number of lines in top margin heading

```

```

225         that appears on each output page; this same number
226         of lines also appears at the bottom margin of
227         each output page.)
228 marMsgType marginMsg(seid cFile);
229         $(Message appearing at top and bottom of each printed
230         physical page, giving owner, security level, etc.)
231 nonDisType nullLevel: $(when die, set maxPrintLevel to nullLevel)
232 VECTOR_OF seid emptyQueue; $(Set queue to emptyQueue when level changes)
233
234     DEFINITIONS
235
236 INTEGER maxNoTextLines IS
237     maxPageLength - 2*marginHeadingLength;
238     $(max number of usable lines per output page.)
239     $(margin message appears at top and bottom of page.)
240
241 INTEGER numPages(textImage file) IS
242     LENGTH(file)/maxNoTextLines +
243     (IF LENGTH(file) MOD maxNoTextLines = 0 THEN 0 ELSE 1);
244     $(Tells how many output pages a file will need to be printed.)
245
246 pageImage bannerPage(line msg) IS
247     VECTOR(FOR i FROM 1 TO maxPageLength:
248         IF i <= marginHeadingLength THEN
249             marginMsg[i]
250         ELSE IF i = marginHeadingLength + 1 THEN
251             msg
252         ELSE IF i INSET {marginHeadingLength+2 .. maxPageLength
253             - marginHeadingLength } THEN
254             " "
255         ELSE marginMsg[i - maxPageLength + marginHeadingLength]
256             $(Put margin message at bottom of page.)
257     )
258     ;
259
260 seid seidAssoc(nonDisType x) IS
261     SOME seid y | TIIinfo(y).nd = x; $(need a seid, rather than nonDisType,
262     to pass as a parameter to SMXflow.)
263
264     EXTERNALREFS
265
266     FROM smx:
267     INTEGER SENlowLevel;
268     INTEGER SENhighLevel;
269     seid: DESIGNATOR;
270     secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
271         tExtent, tNull};
272     privType: {
273         privFileUpdateStatus,     privLink,         privLockSeg,
274         privModifyPriv,           privMount,
275         privSetFileLevel,         privSetSegProcLevel,
276         privStickySeg,           privTerminalLock,
277         privViolSimpSecurity,     privViolStarSecurity,
278         privViolSimpIntegrity,   privViolStarIntegrity,
279         privViolDiscrAccess,     privSignal,     privWalkPTable,
280         privHalt,                privKernelCall, privViolCompartments,

```

```

281     privRealizeExecPermissions};
282
283 daMode: {daRead, daWrite, daExecute};
284 securityCat: DESIGNATOR;
285 integrityCat: DESIGNATOR;
286 domainType: {userDomain, supervisorDomain};
287
288 VFUN TIIinfo(seid s) -> tiiStruct st;
289 VFUN SMXflow(seid pSeid, oSeid; daType da) -> BOOLEAN b;           $(SMXflow)
290 VFUN SMXdap(seid pSeid, oSeid; daType da) -> BOOLEAN b;
291 VFUN SENSEidNsp(seid anySeid) -> INTEGER nsp;
292 VFUN SENSEidToInt(seid anySeid) -> INTEGER i;
293 VFUN SENSEidType(seid s) -> secureEntityType set;                 $(SENSeidType)
294 VFUN SMXhasPriv(seid pSeid; privType privl) -> BOOLEAN b;         $(SMXhasPriv)
295
296     FROM fca:
297 VFUN FCAinfo(seid fSeid) -> globalData gl;                         $(FCAinfo)
298
299     FUNCTIONS
300
301 VFUN spooledFiles() -> SET_OF seid controlFiles;
302     $(Assumed to be set by spooler; for simplicity in specs,
303     will only contain control files, rather than <control file,
304     text file> pairs. A separate VFUN, cToTfile, will give the
305     text file corresponding to any control file in spooledFiles.)
306
307     HIDDEN;
308     INITIALLY controlFiles = {};
309
310 VFUN cToTfile(seid controlFile) -> seid textFile;
311     $(Assumed to be set by spooler; contains the text file associated
312     with a given control file, which is represented in the implementation
313     as a set of <control file, text file> pairs.)
314
315     HIDDEN;
316     INITIALLY textFile = ?;
317
318 VFUN copies(seid cFile) -> INTEGER x;
319     $(Assumed set by spooler. Gets control file info; specifically,
320     how many copies of the file to print.)
321
322     HIDDEN;
323     INITIALLY x = ?;
324
325 VFUN textOfSeid(seid x) -> textImage y;
326
327     HIDDEN;
328     INITIALLY y = ?;
329     $(Assumed set by spooler.)
330     $(An abstraction away from the representation via the directory structure.)
331
332 VFUN currentState() -> dspStates x;
333
334     HIDDEN;
335     INITIALLY x = stopped; $(Requires an interrupt to get started.)
336

```

```

337 VFUN queue() -> VECTOR OF seid qe;
338 $(One of the main VFUNS in despooler. Files are copied from
339 spooledFiles to queue by the OFUN fillQueue. From queue, the
340 files are printed one at a time. Files remain on both spooledFiles and
341 queue until the files are printed, or determined to be unprintable.
342 When printed, they are removed from both spooledFiles and queue.
343 )
344
345 INITIALLY qe =?;
346
347
348 VFUN maxPrintLevel() -> nonDisType x;
349 $(Only files at or below maxPrintLevel may be printed. This level
350 is set, in the implementation, at the beginning of the despooler code.
351 It is set in these specs via the OFUN initPrLevel.
352 It may be changed via the IPC command setLevel.")
353
354 $(In the implementation, there are two separate "filters"
355 appropriate in file movement. active.maxLevel is an upper bound
356 on security level for files moving from the outside world to spooledFiles,
357 and maxPrintLevel is an upper bound on files moving from spooledFiles
358 to queue. In the specs, for simplicity, the same filter, maxPrintLevel,
359 is used in both circumstances.")
360
361
362
363 HIDDEN;
364 INITIALLY x=?;
365
366
367 VFUN output() -> outStream x;
368 $(output is a write-only VFUN simulating actual page printing
369 by appending new pages to the right end of of the output vector.)
370
371
372 INITIALLY LENGTH(x) = 0;
373
374 $(The following VFUNS are the interrupt buffers.)
375
376 VFUN nowbuf() -> BOOLEAN x;
377 INITIALLY x = FALSE;
378
379 VFUN levbuf() -> BOOLEAN x;
380 HIDDEN;
381 INITIALLY x = FALSE;
382
383 VFUN startbuf() -> BOOLEAN x;
384 HIDDEN;
385 INITIALLY x = FALSE;
386
387 VFUN stopbuf() -> BOOLEAN x;
388 HIDDEN;
389 INITIALLY x = FALSE;
390
391 VFUN diebuf() -> BOOLEAN x;
392 HIDDEN;

```

```

393  INITIALLY x = FALSE;
394
395 VFUN newLevbuf() -> nonDisType x;
396  HIDDEN;
397  INITIALLY x = ?;
398
399 $(The following VFUNS correspond to ILPL events.)
400
401 VFUN VstartE() -> BOOLEAN x;
402  INITIALLY x = FALSE;
403
404 VFUN VstopE() -> BOOLEAN x;
405  INITIALLY x = FALSE;
406
407 VFUN VdieE() -> BOOLEAN x;
408  INITIALLY x = FALSE;
409
410
411 $(-----OFUNS-----)
412
413 OFUN processInterrupts();
414  $( "Check interrupt buffers and take appropriate action. Interrupts
415    must be in accordance with above state transition diagram." )
416
417 EXCEPTIONS
418  badTrans1: currentState() = started AND startbuf();
419  badTrans2: currentState() = dead;
420  badTrans3: currentState() = stopped AND
421             (stopbuf() OR diebuf() OR levbuf());
422  $(The above exceptions rule out illegal commands from each state.
423    Handling of multiple legal interrupts from state started is as
424    follows. Set level and stop occurring as multiple interrupts
425    effectively stop the abstract machine at the new level. Set level
426    and die, but not stop, occurring as multiple interrupts essentially
427    die at the null level, with the queue cleared. Stop and die
428    occurring as multiple interrupts effectively treat the interrupts
429    sequentially, first processing the stop. Set level, stop, and
430    die occurring together effectively group set level and stop
431    together as described above, and then sequentially handle the
432    die. "Sequentially" is used here wrt consecutive calls
433    of processInterrupts at the ILPL level.
434
435    This treatment of interrupts is an abstraction and simplification
436    vis-a-vis the implementation, which queues the interrupts on the
437    same queue that files reside, but is logically equivalent to the
438    implementation.
439  )
440
441
442 EFFECTS
443  IF levbuf() AND stopbuf() THEN
444    ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output())+1:
445      IF i <= LENGTH(output()) THEN output()[i]
446      ELSE bannerPage("new level and stop")) AND
447    ^levbuf() = FALSE AND
448    ^currentState() = stopped AND

```

```

449   ^maxPrintLevel() = newLevbuf() AND
450   ^queue() = emptyQueue AND
451   ^VstopE() = TRUE AND
452   ^stopbuf() = FALSE
453
454   ELSE IF levbuf() AND diebuf() THEN
455     ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output())+1:
456       IF i <= LENGTH(output()) THEN output()[i]
457       ELSE bannerPage("new level and die")) AND
458     ^levbuf() = FALSE AND
459     ^currentState() = dead AND
460     ^maxPrintLevel() = nullLevel AND
461     ^queue() = emptyQueue AND
462     ^VdieE() = TRUE AND
463     ^diebuf() = FALSE
464
465   ELSE IF levbuf() THEN
466     ^maxPrintLevel() = newLevbuf() AND
467     ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output())+1:
468       IF i <= LENGTH(output()) THEN output()[i]
469       ELSE bannerPage("new level")
470     ) AND
471     ^levbuf() = FALSE AND
472     ^queue() = emptyQueue
473
474   ELSE IF stopbuf() THEN
475     ^currentState() = stopped AND
476     ^VstopE() = TRUE AND
477     ^stopbuf() = FALSE AND
478     ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output())+1:
479       IF i <= LENGTH(output()) THEN output()[i]
480       ELSE bannerPage("stopped")
481     )
482
483   ELSE IF diebuf() THEN
484     ^currentState() = dead AND
485     ^maxPrintLevel() = nullLevel AND
486     ^VdieE() = TRUE AND
487     ^diebuf() = FALSE AND
488     ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output())+1:
489       IF i <= LENGTH(output()) THEN output()[i]
490       ELSE bannerPage("dead")
491     )
492
493   ELSE IF startbuf() THEN
494     ^currentState() = started AND
495     ^maxPrintLevel() = newLevbuf() AND
496     ^queue() = emptyQueue AND
497     ^VstartE() = TRUE AND
498     ^startbuf() = FALSE AND
499     ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output())+1:
500       IF i <= LENGTH(output()) THEN output()[i]
501       ELSE bannerPage("started")
502     )
503   ELSE TRUE
504   :

```

```

505
506
507 OFUN issueCommand(commands cmd; nonDisType level; BOOLEAN imm);
508   $(Set the appropriate interrupt buffers based on the command.)
509
510   $(****Put in exception check that issueCommand is invoked by
511     someone at or above operator integrity.)
512
513
514 EFFECTS
515   IF cmd = setLevel THEN
516     `levbuf() = TRUE AND
517     `newLevbuf() = level
518   ELSE IF cmd = start THEN
519     `startbuf() = TRUE AND
520     `newLevbuf() = level
521   ELSE IF cmd = stop THEN
522     `stopbuf() = TRUE
523   ELSE IF cmd = die THEN
524     `diebuf() = TRUE
525   ELSE
526     TRUE;
527
528   IF imm THEN
529     `nowbuf() = TRUE
530   ELSE
531     TRUE;
532
533 OFUN spool(seid cFile, tFile: INTEGER ncopies; textImage text);
534   $(Move the file from the "outside world" to spooledFiles. File
535     must be at or below maxPrintLevel. In the implementation,
536     active.maxLevel is used instead of maxPrintLevel in moving files
537     to spooledFiles, and maxPrintLevel is used in moving files from
538     spooledFiles to queue. For simplicity in these specs, the same
539     discriminator, namely maxPrintLevel, is used for both.)
540
541 EXCEPTIONS
542   badLevel: NOT SMXflow(cFile, seidAssoc(maxPrintLevel()), {daWrite});
543
544 EFFECTS
545   `spooledFiles() = spooledFiles() UNION {cFile};
546   `cToFile(cFile) = tFile;
547   `copies(cFile) = ncopies;
548   `textOfSeid(tFile) = text;
549
550
551 OFUN fillQueue();
552   $("Copy files from spooledFiles to queue, without currently removing
553     them from spooledFiles. Only files whose current level is less
554     than or equal maxPrintLevel are copied. ")
555
556 DEFINITIONS
557
558   SET_OF seid queueableFiles IS {seid x | x INSET spooledFiles() AND
559     $(1. Not in queue) (FORALL INTEGER j INSET {1 .. LENGTH(queue())} :
560       queue()[j] ~= x)

```



```

561     AND $(2. ) SMXflow(x, seidAssoc(maxPrintLevel()), {daWrite})
562     }
563     ; $(queueableFiles is the set of files about to move onto queue)
564
565     VECTOR OF seid setToVec(SET OF seid x) IS
566     SOME VECTOR OF seid y | LENGTH(y) = CARDINALITY(x)
567     AND (FORALL INTEGER i INSET {1 .. CARDINALITY(x)}:
568         y[i] INSET x)
569     AND (FORALL seid j INSET x: EXISTS INTEGER k INSET
570         {1 .. CARDINALITY (x)}: j = y[k])
571     ; $(setToVec merely "linearizes" a set into some arbitrary order.
572     In the implementation, the order in which files are queued
573     is based partly on file length. In these specs, the order is arb.)
574
575
576     EXCEPTIONS
577
578     XnoObj: EXISTS seid x INSET spooledFiles():
579         TIIinfo(cToTfile(x)) =?;
580     XnoFile: EXISTS seid x INSET spooledFiles():
581         FCAinfo(cToTfile(x)) =?;
582
583
584     ASSERTIONS
585     LENGTH(queue()) = 0;
586     $(" Whenever fillQueue is invoked the queue is empty because.
587     queue starts out empty; after queue is filled by fillQueue,
588     each element on the queue is processed prior to fillQueue being
589     reinvoked, or else an interrupt occurs forcing queue to be cleared
590     prior to fillQueue being reinvoked.")
591
592
593     EFFECTS
594
595     ^queue() = VECTOR(FOR i FROM 1 TO CARDINALITY(queueableFiles):
596         setToVec(queueableFiles)[i]) ;
597
598
599     OFUN processQelement();
600     $(dequeue the first element of queue and print associated file.)
601
602     DEFINITIONS
603     seid cFileSeid IS queue()[1];
604     textImage text IS textOfSeid(cToTfile(cFileSeid));
605     INTEGER curLength IS numPages(text) + 2; $(2 extra: banner & trailer)
606
607     EXCEPTIONS
608     emptyQ: LENGTH(queue()) = 0;
609
610     EFFECTS
611     ^output() = VECTOR(FOR i FROM 1 TO LENGTH(output()) +
612         copies(cFileSeid)*(numPages(text)+2):
613         $(banner + trailer page
614         for each copy of the file)
615     IF i <= LENGTH(output()) THEN output()[i]
616     ELSE IF (i - LENGTH(output())) MOD curLength INSET {0,1} THEN

```

```

617     bannerPage("normal file")
618     ELSE IF (1 - LENGTH(output())) MOD curLength = curLength -1 THEN
619         $(last page of text may not be full)
620         VECTOR(FOR j FROM 1 TO LENGTH(text) MOD maxNoTextLines:
621             text[LENGTH(text) - (LENGTH(text) MOD maxNoTextLines) +j]
622             $(****Does not yet account for top and bottom margin)
623         )
624     ELSE
625         VECTOR(FOR j FROM 1 TO maxNoTextLines:
626             text[ (((1-LENGTH(output() )) MOD curLength) - 2)
627                 * maxNoTextLines + j]]
628             $(****Does not yet account for top and bottom margin)
629         ) $("end of 'output() = VECTOR...")
630     AND
631     ^spooledFiles() = spooledFiles() DIFF {cFileSeid}
632     AND
633     ^cToFile(cFile) = ?
634     AND
635     ^copies(cFile) = ?
636     AND
637     textOfSeid(cToFile(cFile)) = ?
638
639     ;
640
641     $(The following OFUNS manipulate the VFUNS corresponding to ILPL events.)
642
643     OFUN OstartE();
644     EFFECTS
645         ^VstartE() = FALSE;
646
647     OFUN OstopE();
648     EFFECTS
649         ^VstopE() = FALSE;
650
651     OFUN OdieE();
652     EFFECTS
653         ^VdieE() = FALSE;
654
655
656
657     END_MODULE

```


AD-A111 878

FORD AEROSPACE AND COMMUNICATIONS CORP PALO ALTO CA W-ETC F/G 9/2
K505 COMPUTER PROGRAM DEVELOPMENT SPECIFICATIONS (TYPE B-5). (K-ETC(U)
DEC 80 MDA903-77-C-0333
WDL-YR7932-REV-3 NL

UNCLASSIFIED

4 18 4
24 11 1 16

END
DATE
FILMED
88
DTIC



```
1 MODULE enc            $(password encryption service)
2 MODULE enc            $(password encryption service)
3 $("
4 COPYRIGHT: 1978, by Ford Aerospace and Communications Corp.
5 ADDRESS:    Ford Aerospace and Communications Corp.
6            Western Development Laboratories
7            3939 Fabian Way
8            Palo Alto, California 94303
9            Attention: Software Technology Dept.
10           Mail Stop: V02
11
12 CPCI:  NKSr
13 CPC:
14 MODULE NAME:  enc password encryption service
15 FUNCTION:  encryption of passwords
16 ASSUMPTIONS:  that the Govt. will give us an algorithm
17 HISTORY:
18    AUTHOR:  Tom Berson
19    VERSION:  1.1 of 10/18/78
20    MODULE TYPE: SPECIAL Specifications
21 SPECIAL NOTES:
22 ")
23
24
25 TYPES
26
27 password: VECTOR_OF CHAR ;
28 encryptedPw: {VECTOR_OF CHAR e|LENGTH(e)=ENCepwLength};
29
30
31 PARAMETERS
32
33 INTEGER ENCepwLength;
34
35
36 FUNCTIONS
37
38 VFUN ENcRypt(password pw)->encryptedPw epw:
39    $(definition of the insides of this function to be provided by
40    KSOS customer)
41
42
43 END_MODULE
```

```

1 MODULE ffn      $(file and terminal functions)
2 MODULE ffn      $(file and terminal functions)
3 $("
4 COPYRIGHT: 1978, by Ford Aerospace and Communications Corp.
5 ADDRESS: Ford Aerospace and Communications Corp.
6           Western Development Laboratories
7           3939 Fabian Way
8           Palo Alto, California 94303
9           Attention: Software Technology Dept.
10          Mail Stop: V02
11
12 CPCI: NKSX
13 CPC:
14 MODULE NAME: ffn file and terminal functions
15 FUNCTION: sets the terminal level
16 ASSUMPTIONS: none
17 HISTORY:
18  AUTHOR: Tom Berson
19  VERSION: 1.1 of 10/18/78
20  MODULE TYPE: SPECIAL Specifications
21 SPECIAL NOTES: none
22 ")
23
24
25 TYPES
26
27 terId: {0..255};
28 openDescriptor: {1..PSTmaxOpenDescriptors};
29 openModes: {omStatusUpdate};
30 accessLevel: STRUCT_OF(INTEGER securityLevel,integrityLevel;
31                        SET_OF securityCat sCatSet);
32 tiiStruct: STRUCT_OF(INTEGER securityLevel; SET_OF securityCat sCatSet;
33                     INTEGER integrityLevel; SET_OF integrityCat iCatSet;
34                     INTEGER da; seid owner,group);
35
36
37 DECLARATIONS
38
39 openDescriptor idx,ody;
40
41
42 EXTERNALREFS
43
44     FROM sen:
45     seid: DESIGNATOR;
46
47     FROM smx:
48     securityCat: DESIGNATOR;
49     integrityCat: DESIGNATOR;
50
51     FROM pst:
52     INTEGER PSTmaxOpenDescriptors;
53
54     FROM ker:
55     VFUN K_get_object_level(INTEGER od)->tiiStruct level;
56     OFUN K_set_object_level(INTEGER od; tiiStruct level);

```

```

57 OVFUN K_open(seid oSeid; SET_OF openModes om; INTEGER stCap)
58     ->openDescriptor od;
59
60     FROM alc:
61 VFUN ALCsetPermissible(accessLevel proposed,maximum)->BOOLEAN b;
62
63     FROM tdb:
64 VFUN TDBexists(terId tid)->BOOLEAN b;
65 VFUN TDBlocked(terId tid)->BOOLEAN b;
66 VFUN TDBstandardSeid(terId tid)->seid s;
67 VFUN TDBsecureSeid(terId tid)->seid s;
68 VFUN TDBmaxLevel(terId t)->accessLevel al;
69
70
71 FUNCTIONS
72
73 OFUN FFNsetTerminalLevel(terId tid;accessLevel al;
74     seid newOwner, newGroup; INTEGER newDa);
75
76     DEFINITIONS
77 $(Note that in these definitions odx is not yet bound=>a scope
78 error. However, odx is bound in the scope where the definitions
79 are used.)
80 seid lowner IS IF newOwner=? THEN K_get_object_level(odx).owner
81     ELSE newOwner;
82 seid lgroup IS IF newGroup=? THEN K_get_object_level(odx).group
83     ELSE newGroup;
84 INTEGER lda IS IF newDa=? THEN K_get_object_level(odx).da
85     ELSE newDa;
86
87     EXCEPTIONS
88 NEffnBadTerminal: ~TDBexists(tid);
89 NEffnTerminalLocked: TDBlocked(tid);
90 NEffnBadLevel: ~ALCsetPermissible(al,TDBmaxLevel(tid));
91
92     EFFECTS
93 LET tiiStruct newLevel=STRUCT(al.securityLevel,al.sCatSet,
94     al.integrityLevel,?,
95     lda,lowner,lgroup)
96 $(set the level of the standard path)
97 IN LET odx=EFFECTS_OF K_open(TDBstandardSeid
98     {omStatusUpdate},
99     IN EFFECTS_OF K_set_object_level(odx,newLevel)
100     $(also set the level of the secure path)
101     AND (LET ody=EFFECTS_OF K_open(TDBsecureSeid(tid),
102     {omStatusUpdate},?)
103     IN EFFECTS_OF K_set_object_level(ody,newLevel));
104
105
106 END_MODULE

```

```

1 MODULE gdb          $(group access authentication data base)
2 MODULE gdb          $(group access authentication data base)
3 $("
4 COPYRIGHT: 1978, by Ford Aerospace and Communications Corp.
5 ADDRESS:   Ford Aerospace and Communications Corp.
6           Western Development Laboratories
7           3939 Fabian Way
8           Palo Alto, California 94303
9           Attention: Software Technology Dept.
10          Mail Stop: V02
11
12 CPCI:  NKSR
13 CPC:
14 MODULE NAME  gdb group access authentication data ba
15 FUNCTION:  group access authentication data base
16
17 ASSUMPTIONS:  none
18 HISTORY:
19  AUTHOR:  Tom Berson
20  VERSION: 1.1 of 10/18/78
21  MODULE TYPE: SPECIAL Specifications
22  SPECIAL NOTES:  none
23 ")
24
25
26 TYPES
27
28 groupName: (VECTOR OF CHAR g|LENGTH(g) >= 1);
29 encryptedPw: (VECTOR OF CHAR e|LENGTH(e) = ENCepwLength);
30 userId: seid;
31 groupId: seid;
32 accessLevel: STRUCT OF(INTEGER securityLevel, integrityLevel;
33                        SET_OF securityCat sCatSet);
34
35
36 PARAMETERS
37
38 encryptedPw nullEpw;
39
40
41 EXTERNALREFS
42
43     FROM sen:
44 seid: DESIGNATOR;
45
46     FROM tii:
47 securityCat: DESIGNATOR;
48
49     FROM enc:
50 INTEGER ENCepwLength;
51
52
53 ASSERTIONS
54
55 FORALL groupName x:
56     CARDINALITY({groupId gid|GDBgroupName(gid) = x}) <= 1;

```



```

57 $(In other words, we require unique group names.)
58
59
60 FUNCTIONS
61
62 $(------ group access control info -----)
63
64 VFUN GDBgroupName(groupId gid)->groupName gn;
65
66     INITIALLY gn=?;
67
68
69 VFUN GDBencryptedPw(groupId gid)->encryptedPw e;
70 $(group passwords are optional)
71
72     INITIALLY e=nullEpw;
73
74
75 VFUN GDBadministrator(groupId gid)->userId admin;
76 $(Only the group administrator may modify his group's password)
77
78     INITIALLY admin=?;
79
80
81 VFUN GDBmaxLevel(groupId gid)->accessLevel l;
82
83     INITIALLY l=?;
84
85
86 $(------ db extraction functions -----)
87
88 VFUN GDBvalidGroupName(groupName gn)->BOOLEAN b;
89
90     DERIVATION
91 EXISTS groupId gid : GDBgroupName(gid)=gn;
92
93
94 VFUN GDBgetGroupId(groupName gn)->groupId gid;
95
96     EXCEPTIONS
97 NEgdbBadName. ~GDBvalidGroupName(gn);
98
99     DERIVATION
100 SOME groupId x | GDBgroupName(x)=gn;
101
102 VFUN GDBpwExists(groupId gid)->BOOLEAN b;
103
104     DERIVATION
105 NOT(GDBencryptedPw(gid) = nullEpw);
106
107
108 END_MODULE

```

```

1 MODULE pfn          $(process functions)
2 MODULE pfn          $(process functions)
3 $("
4 COPYRIGHT: 1978, by Ford Aerospace and Communications Corp.
5 ADDRESS:   Ford Aerospace and Communications Corp.
6           Western Development Laboratories
7           3939 Fabian Way
8           Palo Alto, California 94303
9           Attention: Software Technology Dept.
10          Mail Stop: V02
11
12 CPCI:  NKSr
13 CPC:
14 MODULE NAME:  pfn process functions
15 FUNCTION:  process functions:
16           setting process level,
17           getting process level,
18           determination of whether we have an active shell at level L
19
20 ASSUMPTIONS:  none
21 HISTORY:
22   AUTHOR:  Tom Berson
23   VERSION:  1.1 of 10/18/78
24   MODULE TYPE:  SPECIAL Specifications
25 SPECIAL NOTES:  none
26 ")
27
28
29 TYPES
30
31 userId:  seid;
32 groupId:  seid;
33 terId:  {1..255};
34 accessLevel:  STRUCT_OF(INTEGER securityLevel,integrityLevel:
35                        SET_OF securityCat sCatSet);
36 tiiStruct:STRUCT_OF(INTEGER securityLevel;SET_OF securityCat sCatSet;
37                    INTEGER integrityLevel;SET_OF integrityCat iCatSet;
38                    INTEGER da; seid owner,group);
39 activeEntry:STRUCT_OF(terId tid;accessLevel al;seid pSeid);
40
41
42 PARAMETERS
43
44 seid PFNuserStarter;
45
46
47 EXTERNALREFS
48
49   FROM sen:
50   seid: DESIGNATOR;
51
52   FROM smx:
53   securityCat: DESIGNATOR;
54   integrityCat: DESIGNATOR;
55
56   FROM ker:

```

```

57 VFUN K_get_process_level()->ttiStruct level;
58 OFUN K_set_process_level(ttiStruct level);
59 OFUN K_spawn(seid pSeid; VECTOR_OF CHAR p1,p2)->seid newSeid;
60
61     FROM udb:
62 VFUN UDBvalidUser(userId uid)->BOOLEAN b;
63 VFUN UDBinitialShell(userId uid)->VECTOR_OF CHAR s;
64 VFUN UDBloginDirectory(userId uid)->VECTOR_OF CHAR d;
65
66
67 FUNCTIONS
68
69 $(----- process level manipulation functions -----)
70
71 OFUN PFNsetProcessLevel(accessLevel al; seid newOwner, newGroup;
72     INTEGER newDa);
73
74     DEFINITIONS
75 seid lowner IS IF newOwner=? THEN K_get_process_level() owner
76     ELSE newOwner;
77 seid lgroup IS IF newGroup=? THEN K_get_process_level() group
78     ELSE newGroup;
79 INTEGER lda IS IF newDa=? THEN K_get_process_level().da
80     ELSE newDa;
81
82     EFFECTS
83 LET ttiStruct old=K_get_process_level()
84     IN LET ttiStruct new=STRUCT(
85         al.securityLevel, al.sCatSet,
86         al.integrityLevel, old.iCatSet,
87         lda, lowner, lgroup)
88     IN EFFECTS_OF K_set_process_level(new);
89
90 VFUN PFNgetProcessLevel()->STRUCT_OF(accessLevel al; us.rId owner;
91     groupId group)pal;
92
93     DERIVATION
94 LET ttiStruct s=K_get_process_level()
95     IN STRUCT(STRUCT(s.securityLevel, s.integrityLevel, s.sCatSet),
96         s.owner, s.group);
97
98
99 $(----- active shell level memory -----)
100
101 VFUN h_activeSet()->SET_OF activeEntry aes;
102
103     HIDDEN:
104     INITIALLY aes=?;
105
106
107 OFUN PFNpostActive(activeEntry a);
108
109     EXCEPTIONS
110 NEpfnActiveAlready: a INSET h_activeSet();
111
112     EFFECTS

```

```

113 ~h_activeSet()=h_activeSet() UNION {a};
114
115
116 OFUN PFNstrikeAllActiveShells(terId ptid);
117
118     EFFECTS
119 ~h_activeSet()=h_activeSet() DIFF
120     {activeEntry ae|ae INSET h_activeSet()
121      AND ae.tid=ptid};
122
123
124 VFUN PFNqueryActive(terId ptid; accessLevel pal)->seid .activeSeid;
125
126     DERIVATION
127 IF EXISTS activeEntry ae:ae.tid=ptid AND ae.al=pal AND ae INSET h_activeSet()
128     THEN ae.pSeid $(not legal in SPECIAL, must be rebound)
129     ELSE? ; $(NB the ? here is a distinguished return)
130
131
132 $(----- common startup function -----)
133
134 OFUN PFNstartUser(terId tid; userId uid; accessLevel al ;
135
136     EXCEPTIONS
137 NEpfnBadUser: ~UDBvalidUser(uid);
138
139     EFFECTS
140 LET seid startedSeid = EFFECTS_OF K spawn(PFNuserStarter,
141     UDBinitialShell(uid),UDBloginDirectory(uid))
142     IN EFFECTS_OF PFNpostActive(STRUCT(tid,al,startedSeid));
143
144
145 END_MODULE

```

```

1 MODULE sdb          $(system data base)
2 MODULE sdb          $(system data base)
3 S("
4 COPYRIGHT: 1978, by Ford Aerospace and Communications Corp.
5 ADDRESS:   Ford Aerospace and Communications Corp.
6           Western Development Laboratories
7           3939 Fabian Way
8           Palo Alto, California 94303
9           Attention: Software Technology Dept.
10          Mail Stop: V02
11
12 CPCI:  NKSR
13 CPC:
14 MODULE NAME:  sdb  system data base
15 FUNCTION:  system level data base
16
17 ASSUMPTIONS:  none
18 HISTORY:
19   AUTHOR:   Tom Berson
20   VERSION:  1.1 of 10/18/78
21   MODULE TYPE:  SPECIAL Specifications
22 SPECIAL NOTES:  none
23 ")
24
25
26 TYPES
27
28 accessLevel:  STRUCT_OF(INTEGER securityLevel,integrityLevel;
29                        SET_OF securityCat sCatSet);
30
31
32
33 EXTERNALREFS
34
35         FROM smx:
36 securityCat: DESIGNATOR;
37
38
39 FUNCTIONS
40
41 VFUN SDBmaxLevel()->accessLevel 1:
42
43         INITIALLY 1=?;
44
45
46 $(and other functions required to support operator and administrator
47   control functions.)
48 END_MODULE

```

```

1 MODULE tdb               $(terminal data base)
2
3 TYPES
4
5 terId:{0..255};
6 accessLevel: STRUCT_OF(INTEGER securityLevel,integrityLevel;
7                         SET_OF securityCat sCatSet);
8
9
10 PARAMETERS
11
12 $(these contribute to the state of terminals between logout and login)
13 INTEGER TDBdefaultDa;
14 seid TDBdefaultOwner,TDBdefaultGroup;
15
16
17 EXTERNALREFS
18
19         FROM sen:
20 seid: DESIGNATOR;
21
22         FROM tid:
23 securityCat: DESIGNATOR;
24
25
26 FUNCTIONS
27
28 $(----- terminal state information -----)
29
30 VFUN TDBexists(terId tid)->BOOLEAN b;
31
32         INITIALLY b=FALSE;
33
34 $( terminals are brought into existence, i.e. configured, by the
35     use of the system configuration editor. They arrive at reboot.)
36
37
38 VFUN TDBlocked(terId tid)->BOOLEAN b;
39
40         INITIALLY b=FALSE;
41
42 $( Terminals cannot be dynamically removed from configuration.
43     They should be locked by the operator, and TDBexists then
44     edited. They will then be gone at reboot. )
45
46
47 VFUN TDBstandardSeid(terId tid)->seid stdSeid;
48
49         INITIALLY stdSeid=?;
50
51
52 VFUN TDBsecureSeid(terId tid)->seid secSeid;
53
54         INITIALLY secSeid=?;
55
56

```

```
57 VFUN TDBmaxLevel(terId tid)->accessLevel 1;
58
59     INITIALLY l=?;
60
61 VFUN TDBdefaultSpeed(terId tid)->INTEGER speed;
62
63     INITIALLY speed=?;
64
65
66 VFUN TDBdefaultKill(terId tid)->CHAR k;
67
68     INITIALLY k=?;
69
70
71 VFUN TDBdefaultErase(terId tid)->CHAR e;
72
73     INITIALLY e=?;
74
75
76 END_MODULE
```

```

1 MODULE uac      $(user access control functions)
2  $( all functions in this module are subfunctions of the
3    secure server, which is responsible for dialog with the
4    user and for assembling the necessary parameters.)
5
6 TYPES
7
8   userId:seid;
9   groupId:seid;
10  terId:{1..255};
11  userName: {VECTOR_OF CHAR un | LENGTH(un) >= 1};
12  groupName: {VECTOR_OF CHAR gn | LENGTH(gn) >= 1};
13  accessLevel: STRUCT_OF(INTEGER securityLevel,integrityLevel;
14                        SET OF securityCat sCatSet);
15  password: {VECTOR_OF CHAR n | LENGTH(n) >= 1};
16  encryptedPw: {VECTOR_OF CHAR e | LENGTH(e)=ENCepwLength};
17
18
19 EXTERNALREFS
20
21      FROM sen:
22  seid: DESIGNATOR;
23
24      FROM tii:
25  securityCat: DESIGNATOR;
26
27      FROM N_:
28  OFUN N_kill_all(seid terSeid):
29
30      FROM alc:
31  VFUN ALCsetPermissible(accessLevel p,max)->BOOLEAN b;
32  accessLevel ALCgroundLevel;
33  accessLevel ALCloginLevel;
34
35      FROM enc:
36  VFUN ENCCrypt(password p)->encryptedPw epw;
37  INTEGER ENCepwLength;
38
39      FROM sdb:
40  VFUN SDBmaxLevel()->accessLevel mal;
41
42      FROM tdb:
43  userId TDBdefaultOwner;
44  groupId TDBdefaultGroup;
45  INTEGER TDBdefaultDa;
46  VFUN TDBlocked(terId tid)->BOOLEAN b;
47  VFUN TDBstandardSeid(terId tid)->seid stdSeid;
48  VFUN TDBmaxLevel(terId tid)->accessLevel mal;
49
50      FROM udb:
51  VFUN UDBvalidUserName(userName un)->BOOLEAN b;
52  VFUN UDBgetUserId(userName un)->userId uid;
53  VFUN UDBencryptedPw(userId uid)->encryptedPw epw;
54  VFUN UDBdefaultGroup(userId uid)->groupId gid;
55  VFUN UDBmaxLevel(userId uid)->accessLevel mal;
56

```



```

57      FROM gdb:
58 VFUN GDBvalidGroupName(groupName gn)->BOOLEAN b;
59 VFUN GDBgetId(groupName gn)->groupId gid;
60 VFUN GDBpwExists(groupId gid)->BOOLEAN b;
61 VFUN GDBencryptedPw(groupId gid)->encryptedPw epw;
62 VFUN GDBmaxLevel(groupId gid)->accessLevel mal;
63
64      FROM pfn:
65 VFUN PFNgetProcessLevel()->STRUCT_OF(accessLevel al;userId owner;
66                                     groupId group)pal;
67 OFUN PFNsetProcessLevel(accessLevel al;userId newOwner;groupId newGroup;
68                          INTEGER newDa);
69 OFUN PFNstartUser(terId tid; userId uid; accessLevel al);
70 VFUN PFNqueryActive(terId tid;accessLevel al)->seid pSeid;
71 OFUN PFNstrikeAllActiveShells(terId tid);
72
73      FROM ffn:
74 OFUN FFNsetTerminalLevel(terId tid;accessLevel al;userId newOwner;
75                          groupId newGroup;INTEGER newDa);
76 VFUN FFNgetTerminalLevel(terId tid)->STRUCT_OF(accessLevel al;
77                                                  userId owner;groupId group;INTEGER da)tal;
78
79
80 FUNCTIONS
81
82 OFUN UAClogin(terId tid; userName un; password pw);
83
84      DEFINITIONS
85 userId uid IS UDBgetUserId(un);
86
87      EXCEPTIONS
88 NEuacTerLoggedIn: ~(FFNgetTerminalLevel(tid).owner=TDBdefaultOwner);
89 NEuacTerLocked: TDBlocked(tid);
90 NEuacBadLogin1: ~UDBvalidUserName(un);
91 NEuacBadLogin2: ~(UDBencryptedPw(uid)=ENCRYPT(pw));
92
93      EFFECTS
94 EFFECTS_OF FFNsetTerminalLevel(tid,ALCloginLevel,
95                               uid,UDBdefaultGroup(uid),600);
96 EFFECTS_OF PFNsetProcessLevel(ALCloginLevel,uid,UDBdefaultGroup(uid),?);
97 EFFECTS_OF PFNstartUser(tid,uid,ALCloginLevel);
98
99
100
101
102
103 OFUN UAClogout(terId tid);
104
105      EFFECTS
106 EFFECTS_OF N_kill_all(TDBstandardSeid(tid));
107 EFFECTS_OF PFNstrikeAllActiveShells(tid);
108 EFFECTS_OF FFNsetTerminalLevel(tid,ALCgroundLevel,TDBdefaultOwner,
109                               TDBdefaultGroup,TDBdefaultDa);
110 $(EFFECTS_OF K_device_function to stty the terminal back to its
111   default speed, kill, and erase characters)
112 EFFECTS_OF PFNsetProcessLevel(ALCgroundLevel,?,?,?);

```

```

113
114
115
116 OFUN UACchangeGroup(groupName gn; password pw):
117
118     DEFINITIONS
119     groupId gid IS GDBgetGroupId(gn):
120
121     EXCEPTIONS
122     NEuacBadGroup: ~GDBvalidGroupName(gn);
123     NEuacBadPw: ~((GDBpwExists(gid)=>GDBencryptedPw(gid)=ENCrypt(pw));
124     NEuacBadLevel: ~((ALCsetPermissible(PFNgetProcessLevel().al,
125                                     GDBmaxLevel(gid)));
126
127     EFFECTS
128     EFFECTS_OF PFNsetProcessLevel(PFNgetProcessLevel().al,?.gid.?);
129
130
131 OFUN UACchangeAccessLevel(terId tid; accessLevel ral);
132
133     DEFINITIONS
134     userId uid IS PFNgetProcessLevel().owner;
135     groupId gid IS PFNgetProcessLevel().group;
136
137     EXCEPTIONS
138     NEuacBadLevel1: ~((ALCsetPermissible(ral,SDBmaxLevel()));
139     NEuacBadLevel2: ~((ALCsetPermissible(ral,TDBmaxLevel(tid)));
140     NEuacBadLevel3: ~((ALCsetPermissible(ral,UDBmaxLevel(uid)));
141     NEuacBadLevel4: ~((ALCsetPermissible(ral,GDBmaxLevel(gid)));
142
143     EFFECTS
144     EFFECTS_OF PFNsetProcessLevel(ral,?,?,?);
145     EFFECTS_OF PFNsetTerminalLevel(tid,ral,?,?,?);
146     LET seid p=PFNqueryActive(tid,ral)
147     IN IF(p=?S(OR p NOT active))THEN EFFECTS_OF PFNstartUser(tid,uid,ral)
148         ELSE TRUE;
149
150
151 END_MODULE

```

```

1 MODULE udb          $(user access authentication data base)
2
3 TYPES
4
5   userId:seid;
6   groupId:seid;
7   userName: {VECTOR_OF CHAR n|LENGTH(n)>=1};
8   encryptedPw: {VECTOR_OF CHAR e|LENGTH(e)=ENCepwLength};
9   accessLevel: STRUCT_OF(INTEGER securityLevel,integrityLevel;
10                          SET_OF securityCat sCatSet);
11
12
13 EXTERNALREFS
14
15     FROM sen:
16   seid:DESIGNATOR;
17
18     FROM tii:
19   securityCat:DESIGNATOR;
20
21     FROM enc:
22   INTEGER ENCepwLength;
23
24
25 ASSERTIONS
26
27 FORALL userName x:
28   CARDINALITY({userId uid|UDBloginName(uid)=x}) <= 1;
29   $(In other words, unique user names are required)
30
31
32 FUNCTIONS
33
34   $(----- user access control state -----)
35
36 VFUN UDBloginName(userId uid)->userName n;
37
38     INITIALLY n=?;
39
40
41 VFUN UDBencryptedPassword(userId uid)->encryptedPw p;
42
43     INITIALLY p=?;
44
45
46 VFUN UDBdefaultGroup(userId uid)->groupId gid;
47
48     INITIALLY gid=?;
49
50
51 VFUN UDBloginDirectory(userId uid)->VECTOR_OF CHAR dName;
52
53     INITIALLY dName=?;
54
55
56 VFUN UDBinitialShell(userId uid)->VECTOR_OF CHAR sName;

```

```
57
58           INITIALLY sName=?;
59
60
61 VFUN UDBmaxLevel(userId uid)->accessLevel l;
62
63           INITIALLY l=?;
64
65
66 $(----- db extraction functions -----)
67
68 VFUN UDBvalidUserName(userName un)->BOOLEAN b;
69
70           DERIVATION
71 EXISTS userId x:UDBloginName(x)=un;
72
73
74 VFUN UDBgetUserId(userName un)->userId uid;
75
76           EXCEPTIONS
77 NEudbBadName: ~UDBvalidUserName(un);
78
79           DERIVATION
80 SOME userId x|UDBloginName(x) = un;
81
82
83 END_MODULE
```

```

1
2 $(" MODULE          udm.specs (Version 1.5,14:55:05)
3                      Module Name (udm.specs)
4                      File Name (/kr/sunix/specs/nksr/s.udm.specs)
5 CONTENTS            directory manager
6 TYPE                SPECIAL.specifications
7 LAST CHANGED:      11/21/80,14:55:05
8
9 ")
10
11 MODULE udm
12
13 TYPES
14 dirName: VECTOR OF CHAR;
15 path: VECTOR OF dirName;
16 nonDisType: STRUCT OF(
17     INTEGER securityLevel: SET OF securityCat securityCatS;
18     INTEGER integrityLevel: SET OF integrityCat integrityCatS);
19 $(integrityCat is typically the null set)
20 daType: SET OF daMode;
21 modeStruct: STRUCT OF(daType ownerMode, groupMode, allMode);
22 ttiStruct: STRUCT OF(nonDisType nd;
23     modeStruct da; INTEGER owner, group; SET OF privType priv);
24
25
26 PARAMETERS
27 seid ROOTseid $(root of the directory system),
28     sameNspSeid, $(special seid whose nsp component is "sameNsp"
29         to indicate whether or not a file or directory x
30         has been mounted. x has been mounted iff
31         its nsp component is not "sameNsp".)
32     zeroFive, $(special seid whose index field corresponds
33         to the implementation value <0,5> which represents
34         mounted files or directories. This parameter is used
35         by the VFUN derivedDirectory.)
36     slashTmpSeid, $(The seid of the directory /tmp.
37         This seid could be defined as
38         simpPathInterpret(pSeid,ROOTseid,<"tmp">),
39         where simpPathInterpret is obtained from
40         pathInterpret (below) by using directory
41         instead of derivedDirectory")
42     usersTmpSeid; $(The seid of the user's unique tmp directory,
43         which exists under /tmp. usersTmpSeid could be
44         defined as simpPathInterpret(pSeid,ROOTseid,
45         <"tmp","uniqueName">), where simpPathInterpret
46         is obtained from pathInterpret by using
47         directory instead of derivedDirectory.")
48
49
50 DEFINITIONS
51 BOOLEAN mountEquiv(seid s1,s2) IS
52     s1 = s2 OR
53     (SENseidIndex(s1) = SENseidIndex(s2)
54     AND
55     SENseidNsp(s1) = SENseidNsp(sameNspSeid));
56 $(True iff s2 = s1 or s2 has been mounted on s1. mountEquiv

```

```

57  is used in existsPath (below). Basic idea is that IF
58  a sequence of unmounts were to be done, then there
59  would be a "direct" path in the definition of existsPath.")
60
61  BOOLEAN existsPath(seid start,stop) IS
62    (EXISTS INTEGER n | n > 1:
63      (EXISTS VECTOR OF dirName d | LENGTH(d) = n-1
64        AND (FORALL INTEGER j | j INSET 1 .. n-1) :
65          NOT d[j] INSET {".", ".."})) :
66      (EXISTS VECTOR OF seid s | (LENGTH(s) = n
67        AND s[1] = start
68        AND s[n] = stop) :
69        (FORALL INTEGER i | i INSET {1 .. n-1} :
70          mountEquiv(directory(s[i],d[i]),s[i+1]))
71      )));
72
73  seid readDir(seid dirSeid; dirName name) IS
74    IF dirSeid = usersTmpSeid AND name = "." THEN
75      slashTmpSeid
76    ELSE IF dirSeid = slashTmpSeid AND NOT name INSET {".", ".."} THEN
77      usersTmpSeid
78    ELSE
79      dirSeid;
80
81  seid pathInterpret(seid pSeid,dirSeid; path p)
82    $( "This definition, in conjunction with the auxiliary one pil
83    given below, essentially tries to work its way down a path
84    p, starting at node dirSeid,
85    a node at a time, returning as a value the last node
86    encountered. Various exceptions must be met successfully
87    before moving from one node to the next. The notion of a path
88    as used in these specs as an abstraction of the standard Unix
89    notion of a path. For example, the path /a/b/c would be
90    represented simply as the sequence of names a b c.
91    ")
92  IS
93    IF isDirectory(dirSeid) THEN
94      $(Starting node must be directory to continue down path)
95      IF
96        (
97          SMXdap(pSeid,dirSeid,{daRead})
98          OR
99          SMXdap(pSeid,dirSeid,{daExecute})
100         )
101        AND
102        SMXflow(pSeid,dirSeid,{daRead})
103        $(Must have discretionary read or execute permission,
104        and must have mandatory read flow to continue down path)
105      THEN
106        IF LENGTH(p) > 0 THEN
107          pil(pSeid,dirSeid,p) $(Start recursive call to move
108                                down the path)
109        ELSE $(path is of length 0, so return dirSeid)
110          dirSeid
111        ELSE $(flow violated)
112          ?
113      ELSE $(starting node is not a directory)

```

```

113      ?
114      ;
115
116
117      seid pil(seid pSeid, dirSeid; path p)
118      $(An auxilliary definition used by pathInterpret)
119      $(We know at the beginning of each recursive call that
120        dirSeid is a directory to which pSeid has access, and that
121        path p is non-null.)
122      IS
123      IF derivedDirectory(dirSeid, p[1]) ~= ? THEN
124          $(There must be a well-defined
125            next node along the path.)
126      IF      (
127              SMXdap(pSeid, derivedDirectory(dirSeid, p[1]), {daRead})
128              OR
129              SMXdap(pSeid, derivedDirectory(dirSeid, p[1]), {daExecute})
130              )
131          AND
132          SMXflow(pSeid, derivedDirectory(dirSeid, p[1]), {daRead})
133          $(Must have discretionary read or execute permission,
134            and must have mandatory read flow to continue down path)
135      THEN
136      IF LENGTH(p) > 1 THEN
137          IF isDirectory(derivedDirectory(dirSeid, p[1])) THEN
138              $(basic recursion: move one node down the path)
139              pil(pSeid, derivedDirectory(dirSeid, p[1]),
140                VECTOR(FOR i FROM 2 TO LENGTH(p): p[i]))
141          ELSE $(There are at least two more nodes along the path,
142              but the next one is not a directory, so return ?)
143              ?
144          ELSE $(The next node is the last one in the path, so return it)
145              derivedDirectory(dirSeid, p[1])
146          ELSE $(Flow or dap violations encountered)
147              ?
148          ELSE $(There is not a well-defined next node along the path)
149              ?
150      ;
151
152
153      EXTERNALREFS
154
155      FROM smx:
156      seid: DESIGNATOR;
157      secureEntityType: {tFile, tDevice, tTerminal, tProcess, tSegment, tSubtype,
158                        tExtent, tNull};
159      privType: {
160          privFileUpdateStatus,    privLink,          privLockSeg,
161          privModifyPriv,          privMount,
162          privSetFileLevel,        privSetSegProcLevel,
163          privStickySeg,           privTerminalLock,
164          privViolSimpSecurity,     privViolStarSecurity,
165          privViolSimpIntegrity,    privViolStarIntegrity,
166          privViolDiscrAccess,      privSignal,      privWalkPTable,
167          privHalt,                 privKernelCall, privViolCompartments,
168          privRealizeExecPermissions};

```

```

169
170 daMode: {daRead, daWrite, daExecute};
171 securityCat: DESIGNATOR;
172 integrityCat: DESIGNATOR;
173 domainType: {userDomain, supervisorDomain};
174
175 VFUN TIIinfo(seid s) -> tiiStruct st;
176 VFUN SMXflow(seid pSeid, oSeid; daType da) -> BOOLEAN b: $(SMXflow)
177 VFUN SMXdap(seid pSeid, oSeid; daType da) -> BOOLEAN b: $(SMXdap)
178 VFUN SENseidNsp(seid anySeid) -> INTEGER nsp: $(SENseidNsp)
179 VFUN SENseidIndex(seid anySeid) -> INTEGER index;
180 VFUN SENmakeSeid(seid exampleSeid; INTEGER index) -> seid rSeid;
181 VFUN SENseidType(seid s) -> secureEntityType set;
182
183
184 ASSERTIONS
185 $(The following assertions deal with the basic tree structure
186 of directories.)
187
188 $(ASSERTION 1)
189 FORALL seid s | isDirectory(s): directory(s, ".") = s;
190 $(Each directory links to itself via ".")
191
192 $(ASSERTION 2)
193 FORALL seid parentSeid; seid childSeid |
194     (isDirectory(parentSeid)
195     AND isDirectory(childSeid)
196     AND parent(childSeid) = parentSeid):
197     ((EXISTS dirName name: directory(parentSeid, name)
198     = childSeid)
199     AND directory(childSeid, "..") = parentSeid);
200 $(Whenever the ghost variable parent points from child to parent,
201 there is a real link from parent to child via some name, and there
202 is a real link from child to parent via "..")
203
204
205 $(ASSERTION 3)
206 FORALL seid s | isDirectory(s): existsPath(ROOTseid, s);
207 $(Every directory is reachable starting from the root.)
208 $(The only link to the root is via the root itself.)
209
210 $(ASSERTION 4)
211 FORALL seid s1; seid s2 |
212     (isDirectory(s1) AND isDirectory(s2) AND s1 ~= s2):
213     (NOT (existsPath(s1, s2) AND existsPath(s2, s1)))
214 $(No cycles exist between distinct nodes.)
215
216
217
218 FUNCTIONS
219
220 $(-----VFUNS-----)
221
222 VFUN directory(seid dirSeid; dirName name) -> seid ansSeid;
223     HIDDEN;
224     INITIALLY

```



```

225      ansSeid =
226      (IF (dirSeid = ROOTseid AND name INSET {".", ".."}) THEN
227          ROOTseid
228      ELSE
229          ?);
230 $(Initialized by mkDir; updated by mkEntry,rwEntry; finalized by remove.)
231
232 VFUN derivedDirectory(seid dirSeid; dirName name) -> seid ansSeid;
233 $(For certain functions, such as path interpret, derivedDirectory
234   rather than directory is utilized.  directory corresponds to
235   the implementation-visible table, whereas derivedDirectory
236   returns a run-time computable value based on this table.)
237 DEFINITIONS
238
239 seid s IS directory(realDir(dirSeid,name),name);
240
241 EXCEPTIONS
242   badLogic: FALSE; $(Placeholder for exceptions)
243 DERIVATION
244   (IF s ~= ? THEN
245       IF SENSEidNsp(s) = SENSEidNsp(sameNspSeid) THEN
246           SENmakeSeid(realDir(dirSeid,name),SENseidIndex(s))
247       ELSE IF SENSEidType(s) = tFile THEN $(mount processing)
248           SENmakeSeid(s,SENseidIndex(zeroFive))
249       ELSE $(non-file type in directory.  This had better be
250           the last component of a path)
251           s
252       ELSE
253           ?
254   );
255
256
257 VFUN isDirectory(seid dirSeid) -> BOOLEAN b;
258   HIDDEN;
259   INITIALLY
260   b = (dirSeid = ROOTseid);
261 $(Initialized by mkDir; finalized by remove.)
262
263 VFUN parent(seid dirSeid) -> seid s;
264   HIDDEN;
265   INITIALLY
266   s = (IF dirSeid = ROOTseid THEN ROOTseid ELSE ?);
267 $(Ghost variable corresponding to the double-dot entry in each
268 directory.  Initialized by mkDir; finalized by remove.)
269
270
271 $(-----OFUNS-----)
272
273 OFUN mkEntry(seid pdirSeid, eSeid; dirName name1)[seid pSeid];
274 $(Link from (dirSeid,name1) to eSeid.  eSeid must not be
275 a directory - if it is, the OFUN mkDir is used instead.)
276
277 DEFINITIONS
278   seid dirSeid IS realDir(pdirSeid,name1);
279
280 EXCEPTIONS

```

```

281 badFlow: NOT isDirectory(dirSeid) OR NOT SMXflow(pSeid,dirSeid,
282         {daRead,daWrite});
283 badDap: NOT SMXdap(pSeid,dirSeid,{daExecute,daWrite});
284 NotSameFileSystem: SENseidNsp(dirSeid) ~= SENseidNsp(eSeid);
285 badarg: isDirectory(eSeid);
286 nameExists: directory(dirSeid,namel)~= ?;
287 RESOURCE_ERROR;
288
289     EFFECTS
290     ~directory(dirSeid,namel) = eSeid;
291
292
293 OFUN mount(seid pdirSeid, nspSeid: dirName name) [seid pSeid,parentpSeid];
294     $("mount a new extent on directory(dirSeid,name), by replacing
295     its nsp field with that of nspSeid.")
296     $(The object being mounted can only be a certified KSOS file system;
297     this is enforced by the mount CPC)
298     $("The additional implicit parameter parentpSeid is the seid
299     of the parent process of pSeid. It is derivable from pSeid by
300     bringing up additional state information from the kernel specs,
301     but for simplicity here is given in this form. Its use here
302     is in the notMountpriv exception below")
303
304     DEFINITIONS
305     seid dirSeid IS realDir(pdirSeid, name);
306     seid d IS directory(dirSeid,name) ;
307
308     EXCEPTIONS
309     badFlow: NOT isDirectory(dirSeid) OR NOT SMXflow(pSeid,dirSeid,
310         {daRead,daWrite});
311     badDap: NOT SMXdap(pSeid,dirSeid,{daRead,daWrite});
312     badName: directory(dirSeid,name) = ?;
313     notMountpriv: NOT privMount INSET TIIinfo(parentpSeid).priv;
314     cannotMountHere: SENseidNsp(sameNspSeid) ~= SENseidNsp(d) ;
315     notFileType: SENseidType(d) ~= tFile; $(The object upon which we are
316         mounting must be a file, which includes
317         directories)
318
319     EFFECTS
320     ~directory(dirSeid,name) = SENmakeSeid(nspSeid,
321         SENseidIndex(directory(dirSeid,name)));
322
323 OFUN unmount(seid pdirSeid: dirName name) [seid pSeid,parentpSeid];
324     $("Undo the mount operation by restoring the nsp field of
325     directory(dirSeid,name) to the value sameNsp.")
326     $("The additional implicit parameter parentpSeid is the seid
327     of the parent process of pSeid. It is derivable from pSeid by
328     bringing up additional state information from the kernel specs,
329     but for simplicity here is given in this form. Its use here
330     is in the notMountpriv exception below")
331
332     DEFINITIONS
333     seid dirSeid IS realDir(dirSeid,name);
334
335     EXCEPTIONS
336     badFlow: NOT isDirectory(dirSeid) OR NOT SMXflow(pSeid,dirSeid,

```

```

337         {daRead,daWrite});
338 badDap: NOT SMXdap(pSeid,dirSeid,{daRead,daWrite});
339 badName: directory(dirSeid,name) = ?;
340 notMountpriv: NOT privMount INSET TIIinfo(parentpSeid).priv;
341
342     EFFECTS
343     ^directory(dirSeid,name) = SENmakeSeid(sameNspSeid,
344         SENseidIndex(directory(dirSeid,name)));
345
346
347 OFUN mkDir(seid pparentSeid; dirName name1) [seid pSeid];
348     S(Create a new directory as a child of parentSeid, linking from
349     parent to child via name1. The child seid is defined below.
350     Initialize the new directory with dot and double-dot entries, and
351     maintain invariance of tree assertions.)
352
353     DEFINITIONS
354     seid parentSeid IS realDir(pparentSeid,name1);
355     seid childSeid IS SOME seid s |
356         TIIinfo(s) = ? AND SENseidNsp(s) = SENseidNsp(parentSeid);
357
358     EXCEPTIONS
359     badFlow: NOT isDirectory(parentSeid) OR NOT SMXflow(pSeid,parentSeid,
360         {daRead,daWrite});
361     badDap: NOT SMXdap(pSeid,parentSeid,{daExecute,daWrite});
362     badSeid: childSeid = ?;
363     nameExists: directory(parentSeid,name1) ~= ?;
364
365     EFFECTS
366     ^directory(parentSeid,name1) = childSeid;
367     ^TIIinfo(childSeid) = TIIinfo(pSeid);
368     ^directory(childSeid,".") = childSeid;
369     ^directory(childSeid,"..") = parentSeid;
370     ^parent(childSeid) = parentSeid;
371     ^isDirectory(childSeid) = TRUE;
372
373
374 OFUN remove(seid pparentSeid; dirName name1)[seid pSeid];
375     S("Unlink the (parentSeid, name1) entry. If the child is a directory,
376     it must be empty, and pass da checks.")
377
378     DEFINITIONS
379     seid parentSeid IS realDir(pparentSeid,name1);
380     seid aSeid IS directory(parentSeid,name1);
381
382     EXCEPTIONS
383     badFlow: NOT isDirectory(parentSeid) OR NOT SMXflow(pSeid,parentSeid,
384         {daRead,daWrite});
385     badDap: NOT SMXdap(pSeid,parentSeid,{daExecute,daWrite});
386     badName: directory(parentSeid,name1) = ?;
387     badDap2: isDirectory(aSeid) AND NOT SMXdap(pSeid,aSeid,{daRead,daWrite});
388     notEmpty: isDirectory(aSeid) AND (EXISTS dirName name2 |
389         NOT name2 INSET { ".", ".." };
390         directory(aSeid,name2) ~= ?);
391
392     EFFECTS

```

```
393  ^directory(parentSeid, name1) = ?;
394  isDirectory(aSeid) => ^directory(aSeid, ".") = ?
395                      AND ^directory(aSeid, "..") = ?
396                      AND ^isDirectory(aSeid) = FALSE
397                      AND ^parent(aSeid) = ?;
398
399  $(The following OFUNs, piRemove, piMkEntry, piLocate, and piChangeDir,
400  are provided for the KSOS emulator. Their purpose here is simply
401  to let the caller know whether an exception would be returned by
402  a call to remove, mkEntry, locate, and changeDir respectively.
403  Thus, in the OFUNs below, the EFFECTS section
```

DATE
FILMED
8