Copyright © 1991 by the
Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street
New York, NY 10017, USA
All rights reserved as an unpublished work.

This is an unapproved and unpublished IEEE Standards Draft, subject to change. The publication, distribution, or copying of this draft, as well as all derivative works based on this draft, is expressly prohibited except as set forth below.

Permission is hereby granted for IEEE Standards Committee participants to reproduce this document for purposes of IEEE standardization activities only, and subject to the restrictions contained herein.

Permission is hereby also granted for member bodies and technical committees of ISO and IEC to reproduce this document for purposes of developing a national position, subject to the restrictions contained herein.

Permission is hereby also granted to the preceding entities to make limited copies of this document in an electronic form only for the stated activities.

The following restrictions apply to reproducing or transmitting the document in any form: 1) all copies or portions thereof must identify the document's IEEE project number and draft number, and must be accompanied by this entire notice in a prominent location; 2) no portion of this document may be redistributed in any modified or abridged form without the prior approval of the IEEE Standards Department.

Other entities seeking permission to reproduce this document, or any portion thereof, for standardization or other activities, must contact the IEEE Standards Department for the appropriate license.

Use of information contained in this unapproved draft is at your own risk.

IEEE Standards Department
Copyright and Permissions
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA
+1 (908) 562-3800
+1 (908) 562-1571 [FAX]

4.2 System Services

Responsibility: Patricia Oberndorf

4.2.1 Overview and Rationale

This clause describes the system services component of the application platform. It presents a reference model for this component and describes the services provided to application software. Those services are those usually considered as part of an operating system or executive and also those services that may be provided by system level entities such as spoolers and device drivers. Standards, current and emerging, that specify the interface to those system services are also described.

System services are a key component of the application platform and represent the focus of the IEEE effort to produce POSIX base standards. A common set of system services provides support for the portability and the interoperability of application software. While other common services can aid application reuse, system services are those that are common to the largest number of applications.

4.2.2 Scope

System services cover those features that users have come to expect from operating systems or executives. They cover the areas of process management, file management, input/output, memory management, and print spoolers. Because there is a wide variety of platform users, ranging from large general purpose time-shared systems to small time-critical, special-purpose systems, services such as timers and clocks, event management, logical device drivers, and system initialization/reinitialization are included. Services related to distributed systems are also discussed, since application software sees these capabilities through the platform.

4.2.3 Reference Model

This subclause identifies the entities and interfaces specific to the system services of the POSIX OSE. The reference model presented here is consistent with and expands upon the reference model of Section 3. It provides the context for the discussion of System Services in this clause. The basis System Services model is shown in Figure 4-2.

This clause describes the system services portion of the application platform as viewed by a software developer (not necessarily the viewpoint of the end user). This view corresponds to the program design level of abstraction.

The system services API provides the interface between the application software and the system services from the source code point of view. The API defines the program designer's means of accessing the functions, objects, and services of the system.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

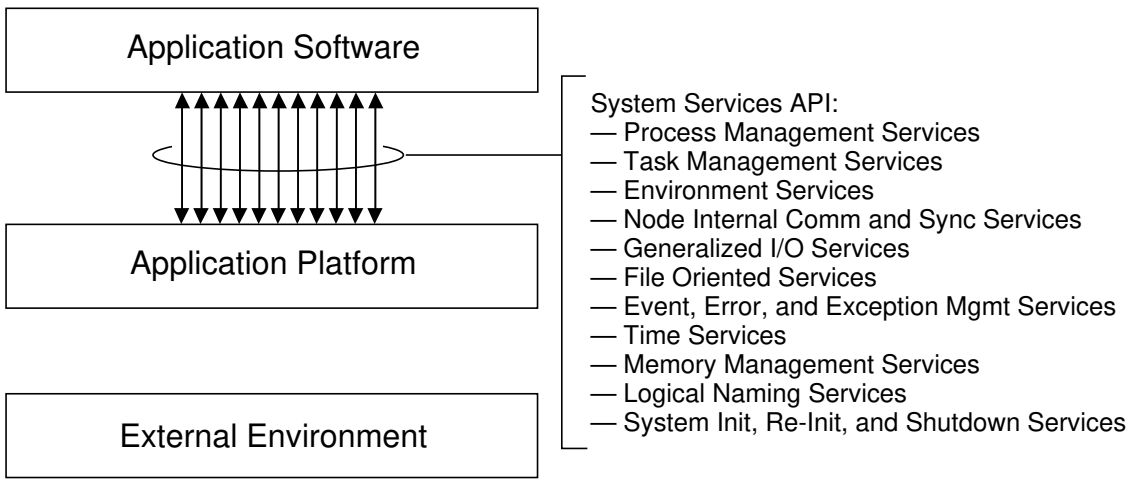

Figure 4-2 – System Services Reference Model

In order for the platform to protect system integrity and ensure system database consistency, application software competing for system resources must access all system resources via system service requests. The formal definition of these requests (or system calls) defines the system services portion of the API.

All of the system services may be available locally or remotely. Some of the system services may be performed remotely if the system is a distributed system with multiple processor nodes. Such distribution is not reflected in Figure 4-2 because it is transparent to users of the System Services.

The platform's device drivers and other software entities are seen as being available to an application program via invocation of the system services. Local devices include sensors, effectors, and connections to independent computing systems. The local devices themselves are a part of the external entities element of the system services reference model. The interfaces used by the application software are the logical device interfaces and are part of the system services. It should be noted that, even though the device drivers are represented within the system services portion of the application platform and the devices themselves are represented within the external entities, there is no unique system service interface illustrated at the EEI in Figure 3-3. This is not an oversight; such interfaces E
are not within the scope of this guide. E

E

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

4.2.4 Service Requirements

This subclause identifies those processor-oriented system services required to support application portability and system interoperability. Subclause 4.2.4.1 describes those system services directly available to an application program via the System Services API. Other processor-oriented services are described in 4.2.4.4. Subclause 4.2.5 identifies the applicable standards.

This subclause describes the major groups of system services that an application may require of a platform. Not all of these services require a programming interface; therefore, services are described as either explicit or implicit services. Explicit services are those that can be accessed from an application program (via the API) and generally are only provided when requested. Implicit services, on the other hand, are services that the platform provides without a direct request. An example of an implicit service is the prevention of one program from writing over the memory of another. An example of an explicit service is a call to a system service routine to output the contents of a block of memory to some device.

4.2.4.1 Application Program Interface Services

This subclause describes the major categories of system services available at the System Services API. These services include:

— Process Management Services

— Task Management Services

— Environment Services

— Node Internal Communication and Synchronization Services

— Generalized Input/Output Services

— File Oriented Services

— Event, Error, and Exception Management Services

— Time Services

— Memory Management Services

— Logical Naming Services

— System Initialization, Reinitialization, and Shutdown Services

4.2.4.1.1 Process Management Services

These services relate to the creation, management, and deletion of processes executing within the scope of an operating system. These processes are distinguished from ‘‘tasks’’ via the following characteristics:

— They have a single thread of execution per address space.

— There is substantial overhead for context switches.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

— Specific attributes are associated only with processes.

In this context, ''management'' consists of those services that affect the execution of a process:

— Stop and restart execution of a process (e.g., suspend, resume)

— Modify processor allocation to a process (e.g., priority, timeslice)

— Modify scheduling of the process based on timer (or other) events

— Protect the process from interruption during critical periods

— Create a process and make it ready for execution

— Destroy a process and recover its resources

— Evaluate a reference to a process

— Evaluate a connection to a process, where a connection is a logical communication path between any two processes

These services schedule or arbitrate the usage of various resources of the OS, particularly the central processing unit (CPU). The scheduling services must be able to queue up requests to use a particular resource. This situation is made more complicated by the common need to schedule processes to run cyclically at a fixed period. When a resource becomes idle, the scheduler must select one of the ''requesters'' of the resource to grant use of the resource. These services are listed separately rather than under the services that use scheduling to emphasize that there should be uniformity and consistency of scheduling across the range of resources.

Typically, there are at least two types of scheduling occurring in an operating system: short-term and long-term. Long-term schedulers determine which possible requesters at a given time may actually request a resource. The short-term scheduler selects from among the active ''requesters'' that currently have need of the resource and allocates the resource to the selected ''requester.'' For example, if the requesters are processes and the resource is the CPU, the long-term scheduler manages the movement of processes from inactive (waiting in batch queues or in hibernation) to active (in wait or execute). The short-term scheduler, on the other hand, would determine which process should execute next on the CPU. Hybrid services between the two may also be available in the operating system.

When a request for a resource is submitted to the operating system (at some local operating system node), it is not always serviced at that local node. The most advantageous way to service the request may result in part or all of the work being performed at a different processor node. Several reasons may cause this to occur, including load balancing, resource availability, computation speedup, hardware preference, and software preference. These services may hide from the application the fact that the functionality was being performed at a different node. This has the advantage that the code needs to know little about the system on which it is running. Alternately, the services may allow the user to specify directly on which logical resource the function should be executed.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The priority scheduling of resources allows the requester to have associated with it its importance to use the service. More complex schemes also have a criticalness of the request that is used for graceful degradation purposes. The scheduler(s) will use the priority information to arbitrate resource requests and to queue requests in the specific order. A priority scheduler may need to support multilevel queues to support proper execution.

Preemptive schedulers will deallocate a resource from a requester when certain events occur. Usually this is when a requester of a higher priority or importance requests the resource or a specified time limit for the resource has expired.

4.2.4.1.2 Task Management Services

These services relate to the creation, management, and deletion of tasks executing within the scope of an operating system. These tasks are distinguished from ‘‘processes’’ via the following characteristics:

— There may be multiple threads of execution per address space.

— There is low overhead for context switches between threads located in the same address space.

In this context, ‘‘management’’ consists of those services that affect the execution of a task:

— Stop and restart execution of a task (e.g., suspend, resume).

— Modify processor allocation to a task (e.g., priority, timeslice).

— Modify scheduling of the task based on timer (or other) events.

— Protect the task from interruption during critical periods.

— Create a task and make it ready for execution.

— Destroy a task.

— Evaluate a reference to a task.

— Evaluate a connection to a task, where a connection is a logical communication path between any two tasks.

4.2.4.1.3 Environment Services

These services provide an application access to a variety of information relating to the operating system environment in which the application is executing. The specific characteristics are:

— Process-specific attributes (process identification, priority, stack size, scheduling attributes, status, memory allocation).

— Task-specific attributes (task identification, priority, scheduling attributes, status, memory allocation).

— Processor-specific attributes (node identification, electronic nameplate information).

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

— User-specific attributes (user identification and terminal ID, user interaction profile).

— Environment variables (command-line arguments, menu selections).

— Current time and date

4.2.4.1.4 Node Internal Communication and Synchronization Services

One or more applications and application subcomponents may run on a processor within an application platform simultaneously. The applications run as independent software entities and communicate among themselves via a variety of mechanisms provided or managed by the system services (see Figure 3-2). An important class of system services relates to the coordination and synchronization of these software entities. In traditional systems, entities execute on a single hardware processor. However, it is becoming common to have multiple processors and networked processors that place more requirements on the system services to provide coordination and synchronization among the many truly concurrent software entities.

When a platform has several software entities executing concurrently, the applications need system services so that the entities can be coordinated and synchronized with each other. With respect to applications written using concurrency, there are two levels of concurrency that are usually seen by the application developer. The first level of concurrency, task level concurrency, is seen when the application is split into multiple subcomponents (tasks) that share access to the data and subprograms of the application. Concurrency services at this level concern the relative priorities and scheduling of tasks within a single application program and their communication with each other. At the second level of concurrency, application level concurrency, a unit is a single application including all its subcomponents. Concurrency services at this level concern the relative importance of the individual applications competing for and sharing system resources.

These services are used to communicate among processes, among tasks, and among processes and tasks residing on the same node. The methods outlined do not include the network specific services described in 4.3, but are limited to methods open to entities executing within the scope of a single operating system. Both synchronous and asynchronous services are defined. The specific services are:

— Create, delete, open, close, read, and write shared memory.

— Create, delete, read, and write event flags.

— Create, delete, set, and wait on semaphores.

— Create/send and receive signals.

— Create, delete, open, close, send to, get from, and control message queues.

— Create, delete, send, and receive streams.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

4.2.4.1.5 Generalized Input/Output Services

These services are used by an application to perform generalized device I/O operations. These operations include synchronous and asynchronous operations for device and class specific functions. Specifically, these form the services needed to implement or include logical device drivers in a system. These services are device initialization, device attachment, asynchronous operation, and error notification. In addition, they include those services that are used to directly access specific device capabilities, particularly those services often referred to as ''raw I/O.''

4.2.4.1.6 File Oriented Services

Mass storage in the form of hierarchy of directories, subdirectories and files will be available to an application executing within the application platform. The following paragraphs describe the services available for creating, accessing, managing, and deleting these entities with mass storage. Both synchronous and asynchronous services are defined.

Naming and Directory Services

These services allow the access of files and directories through logical names rather than the actual hardware device naming conventions. The services allow sharing of files at various levels. For example, the services may not allow any shared naming of files and directories between systems, or they may allow shared files by explicit naming, or they may allow shared files by implicit naming. The directory services present a view or views of the directory structure to the application or target system operator.

File Modification Primitives

Primitive services for files and directories are:

— Read a portion of the file.

— Write to a portion of the file.

— Open access to a file.

— Create a new file.

— Close access to a file.

— Delete a file.

— Copy a file.

— Merge two or more files.

— Append one file to another.

— Split one file into two or more files.

— Support read and write locks at both the record and file levels.

These services may be very complex. For example, the access to read or write may be direct (by record number), sequential (one record at a time), or indexed (by

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

a key). The services must also support a variety of file structures, including linked, segmented, contiguous, serial, and directory.

File Support Services

Additional services support the physical devices on which the files and directory reside. These services include the dismounting/mounting of medium, the formatting of medium, and the partitioning of media.

Realtime Files

Realtime systems often need special files to ensure fast, bounded, and consistent performance in time critical situations. The need for a bounded response time for a given I/O function drives the design of these files and services. One service preallocates the complete disk space needed for a file at creation time, while another guarantees that records within files are aligned in an optimal way (such as along word boundaries). Services support the access of records within the file in ways that make response time constant or bounded, including by direct access.

4.2.4.1.7 Event, Error, and Exception Management Services

These services provide a common facility for the generation and communication of asynchronous events among the system and application programs. A major use of the event services is to report error conditions, but they are also used by device drivers and the platform to provide an indication of some condition to the application programs. These services are:

— Event and error receipt.

— Event and error distribution.

— Event and error management, including user-selectable error processing alternatives (filtering, retry, ignore, accumulate occurrences).

— Event logging.

— Enable/disable and mask/unmask interrupts.

4.2.4.1.8 Time Services

Timers may be a static or dynamic resource on the system, necessitating a variety of allocation and management strategies. These services are used by applications to perform a variety of services based on absolute and relative time. These services are:

— Create a timer.

— Delete a timer.

— Initiate the measurement of an arbitrary specified time duration.

— Receive an indication when the specified duration has elapsed.

— Read the current value of a timer.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

— Initialize a timer with a value and count direction (i.e., increment or decrement).

— Trigger a timer to begin incrementing or decrementing.

— Associate with a timer some action to be taken when the specified duration has elapsed.

4.2.4.1.9 Memory Management Services

These services are used by application processes and tasks to request additional memory and return it to the processor for reuse. They cover the services required to fulfill the needs of both virtual and fixed memory. Specifically, there is a service for locking pages in real memory to support the needs of virtual memory systems.

4.2.4.1.10 Logical Naming Services

These services allow the usage of system resources through logical names rather than the actual hardware device naming conventions. Furthermore, they allow the resources of other processor nodes to be accessed via a logical name so that no knowledge of the resource's location is needed and the resource's location may change over time. Logical names are also used by security services to hide resources from unauthorized processes by only letting authorized processes know the logical name that is needed to use the physical resource.

The logical name to physical name relationship can be one to many, many to one, or many to many. Many times, one physical resource may have multiple logical names as well as one logical name representing a ''bank'' of available physical resources. These services must provide the proper resolution of names, logical and physical, in all of these cases.

4.2.4.1.11 System Initialization, Reinitialization, and Shutdown Services

System initialization consists of services for a complete restarting of the software, starting up the attached hardware subsystems devices, doing subsystem and system self tests, and completely initializing the database.

System reinitialization consists of services for restarting the software while using the existing database information. The software may have to be reloaded and the database may have been reestablished by a system recovery. Attached hardware subsystems may also need to be reinitialized.

Reinitialization also includes a function to restart applications redistributed to other processors after a processor module failure. Within a processor, there is a service to initialize applications in a system with the existing software, but with the database reinitialized. Also within a processor, there is a service to restart the applications in a system with the existing software and database retained.

Shutdown services are those required to perform planned orderly shutdown at the local and remote levels for each and all processor(s) throughout a system. These services support both crisis and non-crisis situations that call for system

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

shutdown. They make sure that the persistent store is in a consistent state, see to the clean termination of all processes, programs, devices, etc., and take care of user notification. They also provide for the running of system diagnostics.

4.2.4.2 External Environment Interface Services

Data Interchange External Environment Interface Services are required by the System Services. Of particular interest are the formats, locations, and procedures for using system administration files, such as password files, system startup files, and configuration files.

4.2.4.3 Interapplication Software Entity Services

This could include support for generalized network/multisession services, such as message handling between system components, global object definition specification, and intermediate language definition.

4.2.4.4 Resource Management Services

These services provide general management functions across the entire platform. They consist primarily of system administration-oriented functions (i.e., management of system interfaces within the scope of the administrator, such as setting up defaults and limits.)

4.2.4.4.1 System Operator Services

The system operator needs to access and control the system services in order to allow the platform to perform properly. If a system has an operator, the major functions that need to be supported are system control, reconfiguration, and status reporting. Currently, these services are usually made available to an operator through a command language interpreter, which is an application program that accesses these system services.

Note that the Windowing Services provide the building blocks (menu utilities, command parsers, etc.) for building the user interface while the System Operator Services make available operating system status and control functions to appropriate application programs with the proper security level.

These services support general conventions and specifications for interaction between system components.

4.2.4.4.2 System Administration

These services and procedures are those required to assure management and allocation of system services to system users, both local and remote. They consist primarily of those services required to establish authorized users of the system, with associated allocation of processor resources, including memory, processor time, priority, and mass storage space. These services are both static (as in the establishment of a new user identification) and dynamic (as in login/logout).

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

4.2.5 Standards, Specifications, and Gaps

Table 4-2 – System Services Standards

Service	Type	Specification	Subclause	E
Process Management	S	ISO/IEC 9945-1	4.2.5.1	E
Task Management	S	ISO/IEC 9945-1	4.2.5.1	E
Environment Services	S	ISO/IEC 9945-1	4.2.5.1	E
Node Internal Comm/Synch	S	ISO/IEC 9945-1	4.2.5.1	E
Generalized I/O	S	ISO/IEC 9945-1	4.2.5.1	E
	G	OSF AES – OSC	4.2.5.3	E
	G	SVID	4.2.5.3	E
File Oriented Services	S	ISO/IEC 9945-1	4.2.5.1	E
Event, Error, and Exception	S	ISO/IEC 9945-1	4.2.5.1	E
	G	OSF AES – OSC	4.2.5.3	E
	G	SVID	4.2.5.3	E
Time Services	S	ISO/IEC 9945-1	4.2.5.1	E
Memory Management	S	ISO/IEC 9945-1	4.2.5.1	E
Logical Naming	S	ISO/IEC 9945-1	4.2.5.1	E
System Init/Reinit/Shutdown	S	ISO/IEC 9945-1	4.2.5.1	E
	G	OSF AES – OSC	4.2.5.3	E
	G	SVID	4.2.5.3	E

4.2.5.1 Current Standards

E

Portable Operating System Interface (POSIX) Part 1

ISO/IEC 9945-1 (IEEE Std 1003.1) is the first in a set of planned international POSIX standards. It defines services and characteristics that need to be in the platform for portable applications, as do some of the other planned standards. Another type of POSIX-related standard is bindings for those services to specific languages. The third type deals with concepts that cross between various groupings of services, such as security and distributed processing.

E

The purpose of the ISO/IEC 9945-1 standard is to define a standard operating system interface based on the UNIX Operating System documentation to support application portability at the source level. The document is intended for systems implementors and applications software developers.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

In addition to ISO/IEC 9945-1, ISO is planning to publish another standard (as yet unnumbered) on test methods for verification of POSIX standards, which will be identical to IEEE Std 1003.3-1991. E E E

E

Table 4-3 outlines the contents of POSIX.1 {2}. This document is identical in its ISO/IEC form (ISO/IEC 9945-1) and the US national standard form (IEEE Std 1003.1). Revisions are currently in progress to deal with: E E

— A language-independent services specification

— A unified data interchange format

— Service interfaces for control of character cell terminals

— Miscellaneous functions identified in comments on the current standard.

Table 4-3 – Functionality of POSIX.1 Standard

File system organization, and file naming conventions
System configuration and file system configuration characteristics
Error messages and reporting mechanism (*errno*)
Application environment information (*environ*)
Process creation, management, and termination: *exec*(), *fork*(), *wait*()
Process environment: user ID, process ID, Group ID
Exception conditions and handling (signals)
Timer operations
File and Directory operations: FIFO files, pipes, status, open/close, read/write
File protection mechanisms
Record and file locking mechanism
Device specific functions: Terminal controls: Processing modes: echo, baud rate, modem termination
C language specific routines: *setlocale*(), nonlocal jumps
User and Group database information (excluding password information)
Data interchange formats (USTAR and CPIO)
Also included is a rationale appendix that provides insight on the selection of various functions and features, including some guidance to developers to understand what types of variations may exist and how that can impact portability.

E

The ISO/IEC 9945-1 standard draws heavily upon major implementations of the UNIX Operating System, including System V and the Berkeley versions. Where a specific behavior was clearly needed (e.g., signals), only a single behavior was permitted. However, there are points where functions were considered optional and

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

others where two different behaviors were considered acceptable. However, in many cases, a solid technical argument favoring one approach over the other was not established. In this case, two behaviors (usually System V and BSD) are defined as being permitted. This is of benefit in writing portable applications, since those that can tolerate both behaviors will run on a wider range of systems. It is also a slight disadvantage in writing such applications, since it can mean handling a wider range of implementations.

NOTE: FIPS 151-1 is a profile of the base standard POSIX.1 {2}. E

E

4.2.5.2 Emerging Standards

E

IEEE P1003.4 E

The IEEE P1003.4 Group is defining realtime extensions to ISO/IEC 9945-1. Draft 9 of the realtime POSIX extensions proposes standardized interfaces to the following functions:

— Response to asynchronous events

— Priority interrupts and scheduling

— Preemptive scheduling

— Memory locking

— High-performance file system (contiguous or other)

— Realtime timers (with nanosecond resolution times)

— Shared memory

— Semaphores

— Interprocess communications (message passing)

— Asynchronous event notification

— Synchronous input and output.

The P1003.4 group is also specifying an interface to threads (P1003.4a).

4.2.5.3 Gaps in Available Standards

While ISO/IEC 9945-1 and P1003.4 both represent very important work, they do not yet address all of the services indicated in 4.2.4. Areas of particular shortfall include Event, Error, and Exception Management Services, some Generalized I/O Services (particularly concerning services for device drivers), and System Initialization, Reinitialization, and Shutdown Services. In addition, Security (see 5.2) and Reliability, Adaptability, and Maintainability services are not reflected in these two base standards, and some capabilities are explicitly considered to be E

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

implementation defined. For some of the services discussed here, adequate consideration is not given to the implications of multiprocessor and distributed implementations of the services and interface provided. Finally, since these are intended to be base standards (or, in the case of P1003.4, an extension to a base standard), profiles are needed in order to select appropriate features and provide appropriate combinations with other related capabilities.

4.2.5.3.1 Public Specifications E

The following are public specifications that define interfaces to services for which no formal standards are currently available.

OSF /1

The Open Software Foundation (OSF) ''Application Environment Specification (AES)—Operating System Component'' (OSC).

Service Gaps Addressed:

— Generalized I/O

— Event, Error, and Exception

— System Init/Reinit/Shutdown

SVID

The AT&T System V Interface Definition (SVID), Issue 3.

Service Gaps Addressed:

— Generalized I/O

— Event, Error, and Exception

— System Init/Reinit/Shutdown

XPG 3 E

X/Open's XPG 3 specifications. E

Service Gaps Addressed: E

— Generalized I/O E

— Event, Error, and Exception E

— System Init/Reinit/Shutdown E

4.2.5.3.2 Unsatisfied Service Requirements

There are two significant areas of the services described above for which no standards currently exist. One is the considerations implied by the use of multiprocessors to implement some or all of the services described herein. The other area is that of interfaces to logical device drivers.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

829 ### 4.2.6 OSE Cross-Category Services

830 #### 4.2.6.1 Capability and Security Services

831 These services support the ability of the system to control usage such that system
832 integrity is protected from inadvertent or malicious misuse. These protection ser-
833 vices provide a mechanism for the enforcement of the policies governing resource
834 usage. Note that many of the security services are implicit services; i.e., they are
835 provided without an explicit request to the operating system. There are two dis-
836 tinct classes of system access with which operating system services must be con-
837 cerned: physical access and logical access.

838 Security services at the physical level are used to protect against security
839 compromise, given unauthorized personnel may have physical access to system
840 hardware. Typically, the physical access is to a terminal and/or terminal/display
841 cables; however, physical access may also include network cables, central process-
842 ing units, disk drives, or tape drives. Prevention of physical access by unauthor-
843 ized personnel may require different operating system services under different
844 circumstances.

845 Logical access is the ability to interact with the operating system via a
846 terminal/display. Security services at the logical level can be implemented
847 through passwords and watchdog timers.

848 Capability services attach operation lists that limit a process's ability to act on
849 resource objects. This is to ensure the resources are not misused. Access to
850 resources can be protected by services using capability lists as well as access lists,
851 lock/key mechanisms, global tables, or through dynamic protection structure ser-
852 vices.

853 **Prevention of Unauthorized Access**

854 The system may need to be guarded from attempted access by unauthorized per-
855 sonnel. The point of access to the operating system that is typically of concern is
856 through the API. Given the mode of operation (system high, multilevel, open) at
857 which the system is operating, these services differ and have differing implica-
858 tions on other system services (such as reliability and naming) and system perfor-
859 mance.

860 **Prevention of Data Compromise**

861 These services prevent access of data by users not authorized to the data. These
862 services may be implemented using access lists on files (and directories) and/or
863 encryption of data or in other ways.

864 **Prevention of Service Denial**

865 These services ensure that a service request will be met by the operating system
866 in a reasonable time if the requester is authorized to use the service. These ser-
867 vices ensure that a bandit user or process cannot cause system malfunction by
868 monopolizing system services or resources.

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

869 **Security Administration**

870 This category involves services to allow the management of the security system,
871 including the administration of permissions to personnel, data, and services as
872 well as capability lists. In addition, it permits the administration access mechan-
873 isms (most often passwords and capability lists) and services that allow the sys-
874 tem to switch modes of operation. The services will likely be accessed by the tar-
875 get system operator with security responsibilities through the target system
876 operator services.

877 E

878 ### 4.2.7 Related Standards

879 The following emerging standards are related to the services covered in this
880 clause, in as much as they address at some level services either explicitly listed in
881 or implied by the services found in 4.2.4:

882 P1003.6 Security Interface for POSIX.

883 P1003.12 Protocol Independent Interfaces (for networks).

884 P1238 OSI Application Program Interfaces (initial effort is to provide at
885 least sufficient facilities for the support of FTAM API
886 specifications).

887 E

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

4.3 Network Services

Responsibility: Charles Severance

4.3.1 Overview and Rationale

This clause describes the network services component of the application platform. It also describes the services provided to application programs and users, and it describes current and emerging standards that are standardizing these services.

Applications gain direct access to network services via the POSIX API. The network is just another system resource (albeit an important one) allocated among the competing processes.

4.3.2 Scope

Network services cover the areas of file transfer, namespace and directory services, electronic mail services, services in support of distributed environments such as remote procedure call, distributed time management, transparent file access, and data representation services. The application programs using these services should be able to access them via a high-level, context-insensitive or low-level, context-dependent interface. E E

In the open systems and distributed system environments, interoperability is of equal or greater importance than portability. The network protocols defined for both Open Systems Interconnect (OSI) and Internet Protocol Suite (IPS) for TCP/IP should provide the basis for the open networking interfaces; however, these interfaces should not preclude the use of some subsequent networking protocol in the future. The interfaces provided by the network services must be network protocol independent and provide for this level of interoperability. E E

It is important for an open system to interoperate with more systems than just other open systems. Many open systems users will have requirements to interoperate with non-OSI networks for the near future. E

4.3.3 Reference Model

This subclause identifies the entities and interfaces specific to the construction of an POSIX Network Environment. This environment is consistent with and extends the environment of Section 3.

As illustrated in Figure 4-3, the components of a network architecture that require standardization are divided into two groups called external environment interfaces (EEI) and application program interfaces (API).

There may be some correspondence between services offered to the application across the API and the interfaces available at the EEI. It is quite possible for an API service to have no corresponding effect at the EEI. A good example of this is an interapplication communication service provided by the Network API between

Copyright © 1991 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.