

ABYSS: A Trusted Architecture for Software Protection

Steve R. White
Liam Comerford

IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Abstract

ABYSS (A Basic Yorktown Security System) is an architecture for the trusted execution of application software. It supports a uniform security service across the range of computing systems. The use of ABYSS discussed in this paper is oriented towards solving the software protection problem, especially in the lower end of the market. Both current and planned software distribution channels are supportable by the architecture, and the system is nearly transparent to legitimate users. A novel use-once authorization mechanism, called a token, is introduced as a solution to the problem of providing authorizations without direct communication. Software vendors may use the system to obtain technical enforcement of virtually any terms and conditions of the sale of their software, including such things as rental software. Software may be transferred between systems, and backed up to guard against loss in case of failure. We discuss the problem of protecting software on these systems, and offer guidelines to its solution. ABYSS is shown to be a general security base, in which many security applications may execute.

Introduction

As computers become a more important source of information and services in our lives, problems of software and data security become increasingly significant. The illicit duplication and use of commercial software is only one example of these problems, but it is increasingly worrisome in the low end of the software market. It constitutes theft of the intellectual property of the authors, in the same way as copyright and patent infringement, and plagiarism. It can be an economic burden as well, since many small software firms rely on sales of a single software package as their only source of income. By disrupting the software marketplace, theft inhibits the growth of this powerful technology.

Since the inception of small computers, attempts have been made to solve the problem of illicit duplication. Technical methods have included writing the application software so that it looks for an unusual, and supposedly uncopyable, feature on the distribution diskette [Voel86], and the attachment of special hardware devices for each application to be used in the system.

These technical methods have not succeeded because of two complementary shortcomings. First, they are not an effective

barrier to duplication. Today's low-end computers are both logically and physically open systems. In fact, most commercial computers of any kind are open, at least to those with operator privileges. The user (or operator) is capable of examining every memory location, and every processor cycle, of the system. Once the behavior of the application is understood, it can be changed to subvert the software protection measures. Similarly, distribution media are completely open to examination and modification. Second, existing technical methods have imposed unacceptable burdens on the legitimate user. Users are often prevented from making backup copies of their software, and from installing their software on hard disks or file servers.

The trend in the software market today is to abandon technical protection methods, and rely solely upon legal protection. This is difficult at best in the widely distributed and anonymous marketplace often addressed by low end software. It is impossible in countries and cultures which have no history of intellectual property law [Chur86]. So, the need for practical technical protection measures continues, and grows.

A practical software protection system must overcome both of the shortcomings outlined above. It must be extremely secure, and ensure that the effort involved in illicitly duplicating an application is at least as large as that involved in rewriting it from scratch. It must also be extremely convenient for the legitimate user, and flexible enough to support a broad spectrum of computing environments and software distribution systems. A variety of authors have explored ideas which go beyond the more common diskette-based protection schemes, in an attempt to meet the requirements of increased security and convenience.

Kent [Kent80] discusses a variety of secure system architectures. He mentions the valuable idea of tamper-resistant modules, which provide physical security, and uses cryptographic techniques to protect applications from exposure. He does not deal with the full spectrum of software distribution methods in use today, nor with the problem of backing up the applications onto another system should the first one fail.

Best [Best79]-[Best84b] presents a crypto-microprocessor approach, in which application software exists in plaintext only within the instruction decoder of the processor. Best's architecture requires the use of a cryptosystem which is not as secure as the modern cryptosystems often used for communications and transactional security. It limits applications

to execution on a single system, and does not deal with the issue of backup. It also requires a substantial restructuring of current computing systems. In some versions of this scheme, a processor per application is required, which is costly.

Arnold and Winkel [Arno85] deal with cryptographic security effectively, but require substantial limitations on distribution methods, the processor, and the execution environment.

Purdy, et al. [Purd82], and Simmons [Simm85], present a protocol for distributing decryption keys for encrypted applications. Their distribution method requires software distributors to maintain a cryptographic facility, which may be a burden in current retail store or mail-order distribution. They also do not deal with the issue of backup.

Everett [Ever85] describes a system similar to the one presented here. Since it requires a public key scheme for key distribution, constantly-updated lists of authorized public keys must be maintained by software vendors. Also, software vendors must maintain cryptographic facilities. No provision is made for backup. No scheme is described which would allow software, installed on one system, to be transferred to another system.

Herzberg and Karim [Herz84], and Herzberg and Pinter [Herz86], present a protection system similar to the one presented here, but they also require software distributors to have a cryptographic facility. They do provide a means of backing up applications in the event of a failure. Their backup method relies on communication with the manufacturer when backed up applications are installed, and on possession of an unused processor on which to install them. This may be problematic in many cases.

Goldreich [Gold86] extends Best's scheme by presenting a method for obscuring the pattern of memory accesses. This increases the protection of Best's scheme, but at the price of additional performance degradation.

Overview of ABYSS

ABYSS (A Basic Yorktown Security System) is an architecture for the trusted execution of application software, and can be used as a uniform security service across the range of computing systems. This paper is oriented towards a solution to the problem of software protection, especially in the lower end of the market. It addresses both security and ease-of-use concerns. Both current and planned software distribution methods are supportable. Users may back up applications at any time, and install them onto any other system in the event of failure, without the intervention of any other party at that time. A general discussion of ABYSS and software protection can be found in [Cina86].

The ABYSS architecture provides the software vendor with tools to enforce the conditions under which the application may be used. Software run under ABYSS executes exactly as it was written, and cannot be modified arbitrarily by the user.

The only information which must be kept secret are certain encryption and decryption keys. Aside from these, all of the

details of both architecture and implementation may be made public without compromising the integrity of the system.

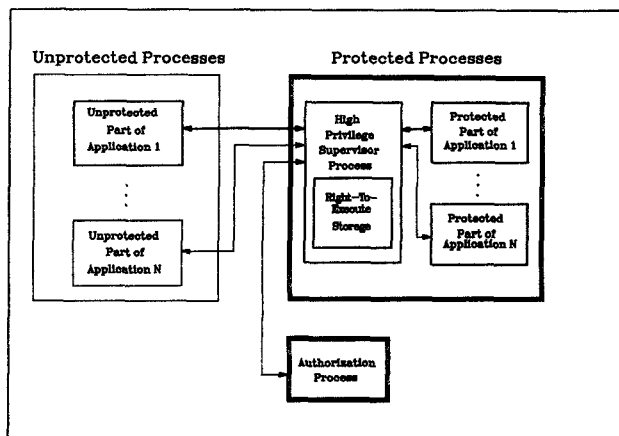


Figure 1. The Architecture of a Protected Processor System

Architecture of ABYSS

The architecture of the system presented here is shown in Figure 1. Applications are *partitioned* into processes which are protected, and processes which are not. Protected application processes are executed within a secure computing environment called a *protected processor*. The conditions under which an application may execute are embodied in a logical object called a *Right-To-Execute*. These conditions are enforced by the protected processor. The movement of Rights-To-Execute into and out of protected processors may require authorization from externally-supplied *authorization processes*.

Protected Processors

A protected processor constitutes a minimal, but complete, computing system. It contains a processor, a real-time clock, a random or pseudo-random number generator, and sufficient memory to store protected parts of applications while they execute. It also contains secure memory for storage of Rights-To-Execute. This storage retains its contents even when the the system power is off.

The protected processor is a logically, physically, and procedurally secure unit. It is logically secure, in that an application cannot directly access the supervisor process, or the protected part of any other application, to violate their protection. It is physically secure (which is indicated by the heavy box in Figure 1), in that it is contained in a tamper-resistant package [Wein86] [Chau84] [Pric86]. It is procedurally secure in that the services which move information, and Rights-To-Execute in particular, into and out of the protected processor cannot be used to subvert the protection.

It is possible for the protected processor to contain the only processors and memory of the entire computing system. Such systems consist of a protected processor, and various peripheral devices such as secondary storage. Alternatively, the protected processor may be part of a larger computing system, and interact with it through the unprotected processes.

In addition to executing protected application processes, the protected processor executes a *supervisor process*. The supervisor process is responsible for ensuring the logical and procedural security of the protected processor. It manages the system's communications resources, to ensure that messages are routed correctly between the protected and unprotected parts of applications. It executes at a higher privilege level than the application processes, and restricts them to isolated protection domains [Denn83] This isolation of application processes from each other, and from their unprotected parts, protects an application process from attacks originating in other application processes, or in the unprotected parts of the computing system.

The supervisor process contains a cryptographic facility for managing encryption/decryption keys. This facility decrypts the protected parts of applications as they are loaded into the protected processor. We place the cryptographic transformation between primary memory (such as RAM) and secondary memory (such as a disk). Best [Best79]-[Best84b] places this transformation between primary memory and the instruction decoder of the processor. Placing it closer to the instruction decoder in the memory hierarchy forces a choice between significant performance degradation of the application, and the use of a cryptosystem with significantly less security than, say, DES.

Placing the transformation between primary and secondary memory, on the other hand, matches the bandwidth of the cryptographic facility to the data transfer bandwidth, allows efficient pipelining of the data to be decrypted, and allows decrypted instructions to be used numerous times without being decrypted each time. It also allows the efficient use of message authentication or manipulation detection codes on parts of the application.

Software Partitioning

For systems in which applications include unprotected processes, it is necessary to partition the application into protected and unprotected parts. The protected part is encrypted when it is outside the protected processor, and only decrypted when it is loaded into the protected processor. The unprotected part is exposed to view.

The protected part executes securely on the protected processor, in that it cannot be examined or modified by any party external to the protected processor. It cannot be examined externally while it is available for execution because of the logical and physical security of the processor, which inhibit all external access to the plaintext of the protected part. The encrypted form of the protected part cannot be modified directly because of the cryptographic techniques used to detect if it has been tampered with. It cannot be modified by rewriting it in a different way, because the partition is chosen so that the protected part is difficult to reconstruct from knowing only the unprotected part.

The partition is designed so that both parts of the application must be present in order to execute the application. Eliminating accesses to the protected part from the unprotected part results

in a nonfunctional application. This intertwining of the protected and unprotected parts extends the logical protection of the protected part to the entire application.

Rights-To-Execute

The software is separated from the right to execute it. Only authorized users of an application have a Right-To-Execute for that application. Rights-To-Execute are created by software vendors, and are used by the supervisor to control the entire range of actions that can be taken with respect to the application.

A Right-To-Execute consists of:

- An encryption and/or decryption key for software packages. This is required to decrypt the application before execution.
- Identifying information about the Right-To-Execute. This can be used to aid the supervisor in searching for the applicable Right-To-Execute.
- Information about how the Right-To-Execute may be used by supervisor software. This allows the software vendor to control what the supervisor may do with the Right-To-Execute. For instance, the software vendor may choose not to allow the Right-To-Execute to be transferred to another protected processor once it is installed.
- Information about how the supervisor may permit the Right-To-Execute to be used by software decrypted under its key. The software vendor may wish to allow the application to change the information in the Right-To-Execute, for instance.
- Information about how the supervisor may permit the Right-To-Execute to be used by non-supervisor software which is not decrypted under its key. It may be useful to allow other applications to report certain information about the Right-To-Execute. For instance, a utility could summarize information about all Rights-To-Execute owned by a user.
- Additional information, at the discretion of software decrypted under the above key. As will be seen later, the application may store such things as an expiration date for its Right-To-Execute, and be assured that the application will not execute after that date.

Authorization Processes

Various supervisor services must be authorized to proceed. For instance, the software vendor must authorize the installation of the Right-To-Execute on a protected processor. Otherwise, it could be installed on any number of protected processors. Clearly, authorization processes must be difficult to forge, to prevent illicit authorizations. Authorization processes may be carried out in a number of ways. Brief descriptions of two of these are given here for clarity in subsequent sections.

Smart Cards: Smart cards are cards the size of a credit card, which contain a microprocessor and memory. They can be constructed to perform a subset of the actions of a protected processor which deal with movement and storage of Rights-To-Execute, but not with application execution. By eliminating the memory in which applications reside during execution, and the real-time clock, current smart cards can contain sufficient cryptographic and storage facilities to authorize supervisor services cryptographically, and to store a limited number of Rights-To-Execute. They can then be used as temporary repositories of Rights-To-Execute being transferred between protected processors, and for a number of other useful services.

Token Cards: Token cards have the same physical appearance as smart cards, but contain a less expensive chip called a token. The token is useful as a one-time-only authorization of supervisor services.

Both smart cards and tokens must be physically secure, to prevent information contained in them from being revealed. Techniques for chip-level security applicable to smart cards and tokens are discussed in [Pric86].

Tokens: Use-Once, Forgery-Resistant Authorizations

We introduce a new authorization mechanism, called a *token process*. The token process is capable of participating in a query-response sequence with a cryptosystem exactly once. Even if the query and response are carried out over insecure channels, the response can still be obtained in such a way that it is extremely improbable that an eavesdropper can forge the behavior of a token process in a subsequent query-response sequence. The token process can be carried out by any simple computing system. It can also be carried out by a small piece of hardware, called a *token*, which is significantly less expensive than hardware capable of cryptographic services such as DES.

How Tokens Operate

Tokens fulfill the following criteria.

- The queries, which are generated randomly by protected processors, are sufficiently numerous that it is extremely improbable that two queries will be the same. Since different queries generate different responses, the response from one query cannot be used as the response to a different query.
- The responses are sufficiently numerous that it is extremely improbable that a random guess at a response will be correct.
- The responses are sufficiently independent of each other, that knowing the response to one query is not significantly helpful in predicting the response to another query.
- The query-response behavior of the token is completely determined by data contained in the token. An encrypted form of these data is delivered to the querying protected processor. This can be done in conjunction with the

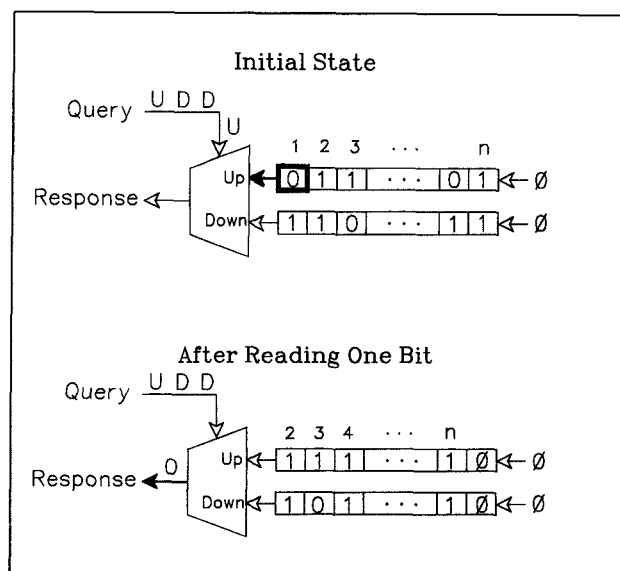


Figure 2. How Tokens Work

query-response sequence, or independent of it. Once the protected processor receives the token data and decrypts it, it can predict the correct response to any query.

- The token data is erased from the token as it is read. This means that a token can only respond to a single query.

Figure 2 shows a simple conceptual realization of tokens. (This is intended to be representational. Real implementations require a small amount of additional support circuitry.) It consists of two shift registers connected to a multiplexor. The registers are shifted left simultaneously in response to a signal on the multiplexor's query line. Each time they are shifted, one bit from either the *up* or the *down* register appears on the output line, depending upon the value of the query bit. At the same time, nulls are shifted into both registers from the right. This cycle is repeated until the token is completely discharged.

The token is loaded by the software vendor by loading random binary strings into both the *up* and the *down* registers. These constitute the token data T_j , and should be effectively unique for each token j . (If an attacker possesses two tokens known to have identical token data, the entire token data can be revealed by querying only the *up* register of one token, and only the *down* register of the other.) The software vendor encrypts this data under a key A , called the *application key*, chosen by the software vendor for a particular application, to form $E_A(T_j)$. The plaintext token data is protected by making the token physically secure against tampering.

The token can then act as a one-time-only authorization from the software vendor, to a protected processor which possesses the application key A . (The means by which the protected processor obtains the application key are discussed later.) To do this, the protected processor reads in and decrypts $E_A(T_j)$ to obtain the token data T_j . It then generates a random query Q , which consists of a string of bits as long as either of the token's registers. The query is presented to the token to obtain the

token's response R . By construction, all of the token data are lost when it is read, even though only half of the data are revealed by the response.

The protected processor can use its knowledge of the complete token data T , to simulate the token, and predict the correct response R' to the query Q . By comparing R to R' , it can determine whether or not the token is a valid authorization, prepared by a party which knows A . Since all of the token data is discharged when it is read, this can only be done once.

Tokens thus serve as counting devices for the software vendor. A vendor who distributes N tokens is guaranteed that there are only N authorizations in the hands of users.

In their ability to prove that they contain certain secret information without revealing a significant fraction of it, tokens resemble the interactive zero-knowledge proofs of [Gold85]. Tokens are much simpler, though, since they only need to respond to a single query about their information.

Forging a Token

Suppose that an attacker has observed the query and response sequence for a token. What is the probability that, armed with this information, the attacker can respond as that token would have to another query by a protected processor? If successful, this would constitute a successful forgery of a token process, and could produce an illicit authorization.

The query to which the attacker must respond is generated randomly, so it will not have a statistically significant correlation to the observed query. The probability of responding correctly to each bit in the query is the probability that that bit in the query is the same as the one previously observed (in which case the attacker knows the correct response), plus the probability that it is different, times the probability of guessing correctly. For a token with n uncorrelated bits in each shift register,

$$p_{\text{forgery}} = [p_{\text{same query}} + (1 - p_{\text{same query}}) p_{\text{correct guess}}]^n \quad (1)$$

If there are no statistically significant correlations present,

$$\begin{aligned} p_{\text{same query}} &= \frac{1}{2} \\ p_{\text{correct guess}} &= \frac{1}{2} \\ \text{so, } p_{\text{forgery}} &= \left(\frac{3}{4}\right)^n \end{aligned} \quad (2)$$

A token with shift registers of length $n = 128$ can be implemented on a very small chip, and gives $p_{\text{forgery}} < 1.02 \times 10^{-16}$.

Since it is a protected processor which generates the query to a token, the protected processor can limit the frequency of queries by controlling the amount of time it takes to generate a query. This inhibits a high-speed "guessing" attack on tokens. The average number of guesses required to come up with a single correct response to a query for a given token is

$$N = - \frac{1}{\log_2(1 - p_{\text{forgery}})} \quad (3)$$

If the time to generate a query is required to be one second, the average time required to forge a response for an $n = 128$ bit token successfully is $t_{\text{average}} > 10^8$ years. Token forgery is covered in more detail in [Stro86].

Secure Sessions Between Protected Processors

Tokens are useful authorization mechanisms when it is impossible or inconvenient to establish a direct communication channel between two protected processors. Where it is possible to establish a such a channel, protected processors possessing an encryption/decryption key in common can set up a cryptographically secure channel between them, mediated by a session key. There are a number of standard ways to do this.

A User's View of the System

The services by which Rights-To-Execute are created, moved, used in executing software, and backed up, are described in detail in the next section. In this section, we illustrate the simplicity of the system to those who use it, by showing a user's view of the system.

For concreteness, we take the case in which applications are purchased through a retail store, as quite a bit of low-end software is today. The applications are intended to be used on a workstation which has an ABYSS architecture, and a hypothetical user interface.

Software Purchase

Retail purchase can be virtually identical to today's practices. The user takes a WonderCalc package off of the shelf, pays for it, and takes it home. No part of the application package need be personalized to the user's system, and the user may remain anonymous. Any application package in the store may be installed on any system, with no prior knowledge of the identity of the system necessary.

Software Installation

The application package contains documentation (hopefully!), a floppy diskette, and a token card. The user inserts the diskette into a drive and the token card into a slot used for just that purpose, and types:

```
install wondercalc
```

The token card authorizes the installation of the right to execute WonderCalc in the user's protected processor. The token is discharged, but the diskette is unchanged. The WonderCalc software may be copied, put onto a hard disk, etc.

Executing Protected Software

From the user's point of view, there is no change in the way in which the application is executed. The user simply types:

```
wondercalc
```

and the application executes as usual.

Backup

The user may create a backup copy of all installed Rights-To-Execute at any time. A smart card is inserted into the slot on the user's system (this is necessary for backup validation, which will be discussed later), and the user types:

```
backup
```

An encrypted file containing the backed-up Rights-To-Execute is written out, to the user's hard disk, for instance, or to a floppy diskette.

If the user's protected processor fails, the backed-up Rights-To-Execute may be installed on another protected processor. This replacement processor may reside in another workstation, or it may simply be a module which is replaced in the user's original workstation. In either case, the installation is accomplished by inserting the smart card into the replacement system's slot, inserting a diskette containing the encrypted file of Rights-To-Execute, and typing:

```
backup -install
```

The Rights-To-Execute are installed on the replacement processor, and the diskette is modified with a message to the hardware manufacturer. The user sends the failed processor, and the diskette, to the hardware manufacturer. The user can use all of the installed Rights-To-Execute as usual.

After a while, the manufacturer sends a diskette to the user, which authorizes the permanent retention of the installed Rights-To-Execute. The user inserts the diskette into the replacement system, and types:

```
backup -retain
```

Transfer

Once installed, Rights-To-Execute are not limited to residing in one system forever. They may be transferred at will by the user, if this is permitted by the software vendor. Suppose that the two protected processors can establish a two-way communication, over a local area network for instance, and that Alice is transferring a Right-To-Execute to Bob. Alice inserts her smart card into her system's slot (this is necessary for backup validation, which will be discussed later), and types:

```
transfer -bob wondercalc
```

The Right-To-Execute is installed in Bob's system, and removed from Alice's.

Supervisor Services

This section describes the supervisor services which were introduced in the previous section. Each of these services can employ either a symmetric cryptosystem or a public key cryptosystem. In this discussion, they are oriented towards a symmetric cryptosystem such as DES, since it is the best match to both the cost and performance requirements of today's lower-end computers.

Each service could use tokens, or communication between protected processors, as authorization mechanisms. Only one of these authorization mechanisms is illustrated in each case for simplicity. We choose the authorization mechanism that is most convenient for illustrating each service. The services are simplified for clarity. An actual implementation would include communication handshaking, verification of the completion of each step of the service, and error recovery.

In order to guarantee the procedural security of the system, each service must ensure four things. It must ensure that no cryptographic keys are exposed in plaintext, that protected software is not exposed in plaintext, that protected software only executes under the terms and conditions chosen by the software vendor, and that the protected software has not been modified.

Creating Protected Software

The ABYSS architecture supports the execution of applications which are not protected, so it can allow existing applications to be migrated to systems with protected processors with no change. To protect an application, the software vendor must create a part of the application to be executed securely, encrypt it, create a corresponding Right-To-Execute, and create an authorization process for installing that Right-To-Execute. The same ABYSS processors which execute protected applications can be used to perform the critical steps in this process, so no special development systems are needed.

Writing Partitioned Applications

The application must be written so that at least part of it resides in the protected processor during execution. Depending upon implementation details, this may be the entire application, or it may be a small fraction of it. The section Secure Software Partitioning discusses the problem of partitioning an application into protected and unprotected parts.

Encrypting the Protected Part of the Application

Once the part P of the application to be protected is complete, it is encrypted under an application key A , chosen by the software vendor, to form $E_A(P)$. The application key may be unique to each application, or even to each copy of each application. This key need not be revealed to anyone else, and the encryption can be done outside of the protected processor, if desired, by the software vendor.

Two goals must be met by this encryption. First, it must avoid the exposure of the plaintext application. Second, it must

prevent the protected part of the application from being modified, even randomly, and executed. This is in support of the requirement that the protected processor execute applications exactly as they were written. If the application could be modified in a known way, it could be instructed to reveal its entire plaintext content. Even if the application could be modified randomly, it is still possible that the random modifications, executing in the protected processor, will cause the revelation of important information about the protected part. The second goal can be met by using message authentication or manipulation detection codes [Juen83] when encrypting the protected part of the application, and having the protected processor use them to authenticate the application before it is executed.

Creating a Right-To-Execute

A plaintext file RTE_A , containing all of the necessary information for the Right-To-Execute for the application, can be created by the software vendor. It contains the application key A , and associated information about how it can be used.

This file is encrypted under a supervisor key S , which is possessed by the protected processor on which the application will be installed, to create $E_S(RTE_A)$. This encryption can be done by a protected processor, as a service to the software vendor. It can be done without revealing the supervisor key S , since the encryption can be done securely inside the protected processor. For simplicity at this point, consider the supervisor key S to be common to all protected processors made by a given manufacturer. Alternatives will be discussed later.

Creating an Authorization Process

The software vendor must now create an authorization process, which will allow the user to install the Right-To-Execute once, but only once. There are many possible authorization processes, including ones in which the user's protected processor must communicate with a processor capable of cryptographic transactions, in order to receive proper authorization. In this section, we will assume that tokens are used as the authorization mechanism.

To prepare tokens as authorizations, the software vendor must generate random token data T_j , which is essentially unique for each token j . The token data are encrypted under the application key to form the encrypted token data $E_A(T_j)$.

Shipping the Protected Software

The above components can now be distributed. The unprotected part of the software U , the protected part $E_A(P)$, the Right-To-Execute $E_S(RTE_A)$, and the encrypted token data $E_A(T_j)$ can be distributed by any means. The token data T_j , which is plaintext, must be kept physically secure, for example in a token.

A typical means of distributing these components would be to package U , $E_A(P)$, and $E_S(RTE_A)$ on a floppy diskette. This diskette could be bulk-reproduced, since every such diskette can be identical. Tokens can be designed so that T_j and $E_A(T_j)$ can

be placed inside of a token. The diskette and token can be packaged together or separately, and shipped to retailers.

Installing a Right-To-Execute

There are several methods by which a user may install Rights-To-Execute on a protected processor. The point of each of these methods is to install the Right-To-Execute for a particular application into the secure, persistent memory of the protected processor, if and only if a valid authorization exists to do so. The use of protected processors in preparing and installing Rights-To-Execute creates a trusted path for distributing Rights-To-Execute. Distribution methods based on communication between two protected processors will be described later. A token-based method proceeds as follows.

1. The user makes the token and distribution diskette available to the protected processor.
2. The encrypted Right-To-Execute $E_S(RTE_A)$ is read into the protected processor, and decrypted under the supervisor key S . This gives the protected processor access to the application key A .
3. The protected processor must now obtain authorization to install RTE_A as follows.
 - a. Read the encrypted token data $E_A(T_j)$ from the token and decrypt it under A to obtain T_j . The fact that T_j and the protected part of the application P are both encrypted under the application key A assures the protected processor that they were created by the same party.
 - b. Generate a random query Q , send it to the token, and obtain a response R from the token.
 - c. The protected processor can now use its knowledge of what the complete token data T_j is supposed to be, and of its query Q , to simulate the token and calculate the response R' expected from a valid token.
 - d. R' is what the token's response to Q should be. If it is different, either the token was invalid, or an attempt was made by an attacker to forge the token's response. There is a valid authorization if and only if $R = R'$.
4. If the authorization was valid, the protected processor installs the Right-To-Execute RTE_A into its secure, persistent memory. If not, RTE_A is purged from the protected processor's working memory and is not installed.

Loading and Executing Protected Software

Once a Right-To-Execute is installed on a protected processor, it may be used at any time to enable the execution of applications associated with it. This is done as follows.

1. The protected processor obtains reference information, which allows it to locate the Right-To-Execute to be used.

2. Once found, the protected processor checks that this Right-To-Execute may be used to execute application programs. Assuming that it is, selected parts of the Right-To-Execute may be made available for reading, and modification, by the application to be loaded.
3. U is loaded into the unprotected memory area, and $E_A(P)$ is loaded into the protected memory area and decrypted. If the decryption of $E_A(P)$ yields a valid message authentication or manipulation detection code, execution of U and P is begun. If not, P is purged from protected memory.
4. At any time during execution, including at the very start, the protected part of the application may be allowed to access (and perhaps modify) selected parts of its own Right-To-Execute. It may use this in conjunction with the real-time clock in the protected processor, for instance, to verify that it is not being executed after its expiration date. Should the protected part of the application find that it is being executed contrary to the terms and conditions specified in its Right-To-Execute, it may effectively terminate its own execution.

Backing Up Rights-To-Execute

If a protected processor fails, the Rights-To-Execute which it contains can become inaccessible. Since they are necessary in order to execute the protected applications to which they belong, the protected applications could become inoperable unless there is some way to back up Rights-To-Execute.

In this section, we describe a method for doing backup which allows the user to install Rights-To-Execute on another processor at any time, without the intervention of the software vendors or the manufacturer of protected processors. A user whose protected processor fails can be up and running on another system immediately. Furthermore, this second protected processor need not be a previously unused one. Another user's protected processor can accept the backed-up Rights-To-Execute, without interfering with any Rights-To-Execute currently installed on it. This backup method nonetheless preserves the integrity of the system, and cannot be used to create extra copies of Rights-To-Execute.

It proceeds in three steps. In the first, a *backup set* of the Rights-To-Execute in a processor is created. This is similar to creating a backup of a hard disk. If the processor fails, the second step is to install the backup set on a second protected processor. The second protected processor can then execute all of the applications that the first one could, in addition to whatever applications it could already execute. The newly-installed Rights-To-Execute have an inactivation date associated with them. They will become temporarily inactive after that date unless a message is received from the hardware manufacturer, authorizing the removal of the inactivation date. The inactivation date may be removed either before or after they have become inactive. After the inactivation date is removed, the Rights-To-Execute behave just as they had in the original processor.

Creating a Backup

There are many ways to perform the backup service. A convenient way is for the user to have a smart card, which is associated with the protected processor to be backed up, and this is the way we will use to illustrate the process. This is not critical, as any other protected processor could be used instead. A single smart card is used here for concreteness. By employing secret-sharing methods [Sham79], multiple smart cards could also be used, to further decrease the effect of a single failure.

The smart card, and its associated processor, each have a supervisor key S_u within them, which is unique to the two of them. That is, the smart card and associated protected processor both possess S_u , but no other smart card or protected processor possesses it. This key can be used to identify them to each other, and to establish a secure session between them with the protocol described in a previous section. Given this supervisor key, a backup set is created as follows.

1. The protected processor and smart card use the unique supervisor key S_u to establish a secure session between each other, which is mediated by a session key B .
2. The collection $C = \{RTE_{A1}, RTE_{A2}, \dots, RTE_{An}\}$ of Rights-To-Execute to be backed up is encrypted under B to form $E_B(C)$.
3. $E_B(C)$ is written to an external storage medium. If there is sufficient storage on the smart card, it may be kept there. It could also be kept on a diskette. This file may be copied as many times as is desired.
4. The smart card stores the key B , and identifies it as the current backup key.

Installing a Backup

Each smart card also contains a supervisor key S_c which is common to all protected processors made by a given manufacturer. If the protected processor fails, this key may be used to establish a secure session between the smart card and any other protected processor.

Protected processors also contain common supervisor keys S_m and S'_m , which are used only for communicating with the hardware manufacturer. S_m is used only for sending information to the manufacturer, and S'_m is used only for receiving information from the manufacturer.

Installation of the backup is done as follows.

1. The smart card and the second protected processor use the common supervisor key S_c to establish a secure session between each other, mediated by a session key \tilde{B} .
2. The backup key B , and the unique supervisor key S_u , are encrypted under \tilde{B} , transmitted to the second protected processor, and decrypted. Remember that S_u is the unique supervisor key of the *first* (failed) processor.

3. The backup key B , and the unique supervisor key S_u , are erased from the smart card. This prevents the smart card from installing the backup set in another protected processor. This also prevents any subsequent backup set from the first processor from being installed.
4. The collection of backed up Rights-To-Execute $E_b(C)$ is read into the second protected processor, decrypted, and installed.
5. The Rights-To-Execute in this collection are made conditional on the existence of a Right-To-Execute associated with S_u , and RTE_{S_u} is given an inactivation date.
6. The second protected processor prepares a message for the hardware manufacturer, $E_{S_m}(S_u)$, and writes it to an external medium. The user sends this medium, along with the first (failed) processor, back to the hardware manufacturer.

Retaining Backed-Up Rights-To-Execute

The Rights-To-Execute in the second processor will continue to operate as usual. If no authorizing message were received prior to the inactivation date, the Rights-To-Execute which were part of the backup process would become inactive, but would not be erased. They could not be used, but would remain in the protected processor, awaiting receipt of the message from the hardware manufacturer.

The hardware manufacturer, upon receiving the first (failed) processor and $E_{S_m}(S_u)$, decrypts the message, and verifies that the serial number of the failed processor corresponds to the one associated with the unique supervisor key S_u . This correspondence is checked to ensure that the correct processor has been returned. If a different processor were returned, and the authorization were given to remove the inactivation date from the second processor anyway, the installed Rights-To-Execute would have been permanently duplicated.

Assuming that the correspondence is verified, the inactivation date is removed as follows.

1. The hardware manufacturer encrypts a message $E_{S_m'}(S_u)$, authorizing the removal of RTE_{S_u} , and the elimination of any dependency of other Rights-To-Execute on RTE_{S_u} . $E_{S_m'}(S_u)$ is sent to the user.
2. The second protected processor reads $E_{S_m'}(S_u)$ and decrypts it. Since it was encrypted under S_m' , it is guaranteed to have originated from the hardware manufacturer. The protected processor then erases RTE_{S_u} , and eliminates all dependencies on it. The installed Rights-To-Execute may now be used just as they were on the original processor.

Transferring Rights-To-Execute

A Right-To-Execute need not be permanently associated with the protected processor on which it was first installed. Conceptually, it is straightforward to transfer a Right-To-Execute From Alice's protected processor to Bob's by establishing a secure session, transmitting the Right-To-Execute from the Alice's processor to Bob's, installing the Right-To-Execute on Alice's processor, and erasing it from Bob's.

A Necessary Connection Between Transfer and Backup

There is a hole in the transfer method outlined above, which permits Rights-To-Execute to be duplicated. This same hole exists in any system which supports both transfer and backup.

Rights-To-Execute could be duplicated as follows.

1. A backup of $C = \{RTE_{A1}, RTE_{A2}, \dots, RTE_{An}\}$ is made from Alice's protected processor.
2. The Rights-To-Execute in C are transferred to Bob's processor. Naturally, they are erased from Alice's processor in the process.
3. The backed-up Rights-To-Execute are installed in Carol's processor. This is unaffected by the fact that Alice's processor (the one backed up) no longer contains these Rights-To-Execute, as a result of transferring them to Bob.
4. Alice's processor is broken, and returned to the manufacturer as failed.
5. The manufacturer verifies the failure, and permits C to be permanently installed in Carol's processor.

At this point, both Bob's and Carol's processors have C installed permanently, and the Rights-To-Execute have been successfully duplicated. While Alice's processor must be sacrificed in the process, this may be a small price to pay if the collection of Rights-To-Execute was very valuable.

Transfer With Backup

This hole is plugged by coupling the services of transfer and backup. Each time a transfer is made, all previous backup sets must be invalidated, and a new one created. A revised transfer service, which transfers RTE_{Ak} from Alice's processor to Bob's, is as follows.

1. Alice's protected processor and its associated smart card use their unique supervisor key S_u to establish a secure session with each other, mediated by the session key B' .
2. $C' = \{RTE_{A1}, \dots, RTE_{Ak-1}, RTE_{Ak+1}, \dots, RTE_{An}\}$ the revised backup set, is encrypted under B' , and stored as described previously, and the associated backup key B' replaces the previous backup key B on Alice's smart card. Alice's smart card is now capable of installing only C' ,

the most recent backup set. It is not capable of installing C , or any other previous backup set.

3. Alice's and Bob's protected processors use a common supervisor key S_c to establish a secure session with each other, mediated by the session key T .
4. $RTE_{A,k}$ is encrypted under T by Alice's processor, transmitted to Bob's, and erased from Alice's.
5. Bob's processor decrypts $RTE_{A,k}$ and installs it.

The method of changing the encryption key for data sets, to ensure that they are up to date, is described in a somewhat different context in [Kent80]. In general, this method permits information, such as Rights-To-Execute to be encrypted, stored on secondary media, and still be guaranteed to be up to date.

Secure Software Partitioning

Software partitioning is necessary if part of the application must execute from unprotected memory. Only that fraction of an application which executes from protected memory is hidden from view. Cost constraints may require that the protected processor be limited in available processing power and memory. Software partitioning is a means of creating a dependence of the untrusted part of the system on the trusted part. This can extend the protection of the trusted part to the entire application.

It is difficult to construct a partitioning method that is both demonstrably secure, and convenient enough to use in practical application development efforts. In this section, we outline some ideas that may lead to a better understanding of this problem.

If it were possible to reconstruct the full application without seeing the protected part, the protection of the application would be compromised. The goal of secure software partitioning is to make this reconstruction prohibitively complex. For software protection, it is sufficient if reconstructing the software is more difficult than rewriting it from scratch. We define two kinds of complexity in partitioned software: semantic complexity and combinatorial complexity.

Semantic Complexity

Semantic complexity reflects the difficulty of reconstructing the protected part by examining the environment of its interaction with the unprotected part. At one end of a spectrum of partitioning methods, selected obscure parts of the application are protected. For instance, an application may contain a proprietary algorithm, all of which could be protected. If that part of the program was difficult to write initially, it is likely to be difficult for an attacker to reconstruct it. At the other end of this spectrum, random

parts of the application could execute in the protected processor. For instance, every tenth line of a program could be executed by the protected processor. This is semantically complex to the extent that it is difficult to understand a program that has a large number of lines missing.

Combinatorial Complexity

Combinatorial complexity reflects the difficulty of exhaustively characterizing the behavior of the protected part by watching what it does. Consider an application in which there are n access points in the unprotected part, at which accesses are made to the protected part. At each access point, a k bit argument is passed to the protected part, and the protected part performs some calculation.

If this results in $\Omega(2^k)$ independent states of the system, essentially all possible values of the argument must be tried by an attacker to completely characterize the effects of that calculation. This is the case, for instance, in a 1-1 function of the argument, whose value is returned by the calculation. This is already daunting if the resulting values are returned to the unprotected part immediately, since it requires the attacker to write a program which executes the protected part $\Omega(2^k)$ times to characterize each access point. The characterization can be made even more difficult if some or all of the results are stored in the protected part instead.

If no state is stored in the protected part, each access made to it by the unprotected part is independent of every other access. In this case, the combinatorial complexity, as measured by the number of times the attacker must execute the protected part, is $\Omega(n)$ in n .

If state can be stored, each access point can write its argument into some location in protected storage, and return a value read from another location. If the identity of these locations are obscured in the unprotected part, by hashing for instance, it is straightforward to show that $\Omega(n^2)$ executions of the protected part are required for characterization. Ideally, a partitioning method should have a combinatorial complexity which is $\Omega(2^n)$, as well as $\Omega(2^k)$.

Partitions can be created which are complex in both senses. Portions of the application which are semantically complex can be executed by the protected processor as independent routines. If each of these routines performs at least one of a set of functions which can be shown to be combinatorially complex, and if these functions are semantically unrelated to the semantically complex part of the routines, then the routines embody both kinds of complexity.

Attacks on the System

Attempts at forging a token process were discussed earlier. In this section, we discuss a number of other possible attacks on the system, and their consequences.

Plaintext Software: The Asset Being Protected

Software protection is an unusual application of cryptography. Often, messages being encrypted have value for only a small amount of time. The contents of a telephone call made a year ago may have little value, for instance. Furthermore, many cryptographic systems can be constructed so that compromise of a single key, or the compromise of a single cryptographic facility, are of limited value.

These things are not true of practical software protection systems. The asset being protected is the plaintext application itself. If an attacker can obtain the plaintext application, the attacker is in the same position as the original software vendor. The application can be copied, altered, and redistributed at will. It may even be reprotected with the protection system before distribution.

Once an attacker obtains the plaintext application, the application's protection is reduced to what exists on unprotected systems today. The application can be made obscure, made to check for odd features on a diskette, and so on, but these offer poor technical protection. Copyright laws, where they exist, provide some degree of legal protection. It may be possible to enhance the legal protection against redistribution of a modified version of the application through the use of fingerprinting techniques [Blak85].

The essential point, which is so often overlooked in software protection discussions, is that, once an attacker obtains the plaintext of the application, the technical protection of that application has been rendered useless.

Cryptanalytic Attacks

The cryptographic techniques used in ABYSS are entirely conventional. Cryptanalytic attacks on keys, or on the plaintext application itself, may be inhibited in standard ways. The cryptographic protection of keys and applications is at least as strong as the protection of the cryptosystem against a non-chosen plaintext attack. Modern cryptosystems, such as DES, are sufficiently resistant to these attacks.

A cryptanalytic attack on an application key could be attempted, using the response from a token query and the encrypted token data. This attack is more difficult than a non-chosen plaintext attack, since not all of the corresponding plaintext is available in the token response.

Physical Attacks

A physical attack on the protected processor itself, if successful, could reveal both supervisor and application keys. But, just as important, it could reveal the plaintext for any application that can be executed on the compromised processor. If the processor is compromised, the attacker can simply purchase every application which is to be compromised. They can each be installed, loaded onto the compromised processor, and their plaintext can be read out in turn.

Physically compromising any processor compromises all software that can be executed on that processor. Since it is necessary for virtually every application to be installable on every processor, the physical security of the protected processor is extremely important. [Wein86] describes a low-cost tamper-resistant package which is useful for protecting information in multi-chip systems. The physical security of smart cards and tokens is also critical, though losses due to a compromised token can be limited to the ability to install a single application. Chip-level security, which is applicable to smart cards and tokens, is discussed in [Pric86].

On Trusting the Hardware Manufacturer

All software protection methods place a great deal of trust in the party that manufactures the protection method. Cryptographic keys must be kept secret, for instance, and must not be used illicitly to reveal information which is encrypted with them.

In [Herz86], the authors maintain that public key cryptosystems are superior to symmetric cryptosystems for software protection, since public key cryptosystems can prevent the hardware manufacturer from knowing the private keys of each protected system. While this is true, it is of limited utility if the hardware manufacturer is dishonest.

The hardware manufacturer could build a "protected" system which is identical to those distributed commercially, but which is not physically secure. Applications could be purchased for that processor as usual, and their plaintext exposed. Although the hardware manufacturer cannot compromise another user's processor *per se*, the applications themselves can be compromised. If the hardware manufacturer cannot be trusted, significant assets are at risk, regardless of what cryptosystem is used.

ABYSS allows each software vendor to distribute only to protected processors made by manufacturers that they trust. This is done by creating Rights-To-Execute which can be loaded only onto processors which have the trusted manufacturer's supervisor key.

Distribution Alternatives

Supervisor keys may be unique to a given protected processor, common to all processors made by a given manufacturer, or anything in between. The two extremes, unique keys and common keys, are discussed below.

Most low-end software distributed today can be run on any system. It is not necessary to know on which system it will be run at the time of purchase, nor to register the system with some central authority before it can be run. We call this *untargeted* distribution, since it is to any member of the entire set of protected processors produced by that manufacturer.

These characteristics can be preserved if common supervisor keys are built into protected processors. The software vendor's protected processor encrypts the

application's Right-To-Execute under a common supervisor key. The Right-To-Execute is then distributed to retail stores. Any user can buy the package, and install it on any system, since any system can decrypt the Right-To-Execute. The software vendor need not be involved in the installation.

Common supervisor keys also allow secure two-way communication between any two protected processors. This is advantageous for transferring and backing up Rights-To-Execute.

Some software packages, particularly expensive software systems, are distributed to only a select group of users. Often, the software vendor wants to have a close contractual relationship with the users, or wishes to maintain a detailed list of users for maintenance purposes. We call this *targeted* distribution, since it is to a select group of systems, known in advance by the software vendor.

Targeted distribution can be achieved with unique supervisor keys. The software vendor uses the unique supervisor key of the target system to encrypt the Right-To-Execute destined for that system. This ensures that no other system will be able to install the Right-To-Execute. The software vendor can obtain the user's unique supervisor key, without it being revealed. This can be done by having the user's protected processor encrypt the unique key S_u under a common supervisor key S_c to form $E_{S_c}(S_u)$. $E_{S_c}(S_u)$ is then sent to the software vendor's protected processor and decrypted.

Unique supervisor keys allow unique identification of any protected processor. This is useful for such things as coupling a smart card to a particular protected processor in the backup service.

New Capabilities

ABYSS opens up a wide variety of new capabilities in computing systems, which were either too expensive, or too inconvenient previously. Two important capabilities are the technical enforcement of the terms and conditions of software sales, and the protection of software distribution channels.

Technical Enforcement of Terms and Conditions

A software vendor can specify the conditions under which the Right-To-Execute can be used for an application, or even for a particular copy of an application. This enables each software vendor to enforce a very wide range of conditions on the use of software, and to do so technically. In many cases, only legal enforcement was available previously.

The ability of supervisor services to use the Right-To-Execute can be limited. For instance, the software vendor may choose not to allow a Right-To-Execute to be transferred to another system, once it is installed. This provides technical enforcement of the condition that execution only occur on a given system.

The software vendor can specify terms and conditions under which the application will execute, and allow the application itself to enforce them. In general, any terms and conditions that can be embodied in data in the protected processor, as a part of the Right-To-Execute, can be enforced. The application has access to these data, and may test the specified conditions against such system facilities as the real-time clock in order to determine if it is being properly used. If it is not, the application simply refuses to execute.

By storing an expiration time and date, or a maximum number of allowed uses of the application, it is possible to enforce the concept of limited lifetime software.

Rental software currently poses a significant problem to software vendors. If the rented application's Right-To-Execute contains a terminal date, and if the application checks this against the real-time clock whenever it is run, the rental agreement is enforced. After the end of the rental period, both the renting company and the software vendor are assured that the renter can no longer use the application.

In an environment where a large number of workstations are connected to a network, such as a local area network with a file server, significant economies can be achieved for widely used applications. A single copy of an application can be kept on a file server. The protected processor associated with this server can store a Right-To-Execute which enables up to some fixed number of copies to execute, for instance. When a user requests the application, a single Right-To-Execute is also transferred to the workstation, and the server's count is decremented. The transferred Right-To-Execute may require renewal every minute or so to continue to be valid. The workstation can request renewal of the Right-To-Execute from the server at any time, and thus continue to possess a valid Right-To-Execute as long as is necessary. If the workstation does not check back with the server, the server knows that the workstation's Right-To-Execute has expired, so the server can increment its count of Rights-To-Execute for that application. This *software lending library* is much like a lending library for books, in which the books automatically reappear in the library when they are due.

Currently, retail stores seldom allow a user to return purchased software, even if the software turns out to be unfit for the user's purpose. In part, this is because the store has no way of verifying that the user has not made an illicit copy of the software, thus depriving the store of revenue. Since a Right-To-Execute can be securely transferred back to the store, ABYSS systems permit such verification, and enable returnable software.

As an extension of this, application packages may be prepared for demonstration purposes, and be sold for a price lower than the normal application. A user can try the full application, not some subset of it, in the particular computing environment, and with the particular data, on which the application is to be used. If the demonstration package is guaranteed to expire after a fixed time, or after some number of uses, the software vendor can use demonstration software as a marketing aid, without losing sales of the application.

Protection of Distribution Channels

Separating the software from its Right-To-Execute permits a large variety of distribution channels to be used for either. Many electronic distribution methods are not employed today because of the difficulty of preventing illicit interception of the applications. ABYSS can protect all of these distribution channels.

In environments in which there can be secure two-way communication between the software vendor and the user, Rights-To-Execute may be distributed on demand to individual workstations. Both software, and Rights-To-Execute can then be distributed on local or wide area networks, or by download from host systems to workstations.

Optical disks and compact disks have been suggested as possible distribution media for software. Ironically, one of their disadvantages is that they are capable of storing so many applications on a single disk, that it is difficult to market the entire set of applications together. With ABYSS, it is possible to sell a disk containing hundreds of protected applications at a very low price, and sell individual Rights-To-Execute separately. Similarly, it is possible to transmit a large collection of applications on FM radio bands, or on a cable television network, and sell their Rights-To-Execute separately. This enables very inexpensive distribution channels for software applications, while retaining their protection.

Additional Security Applications

An application which executes in a protected processor does so in exactly the manner specified by the software vendor. Attackers cannot alter the terms and conditions under which it executes, or its actions when it is executing. In security applications, the protected processor serves as a trustworthy agent of the supplier of computing services, much as it does for the software vendor. This makes protected processors an ideal base for many security applications.

Given an appropriate implementation of the protected processor, security applications such as dial-in host access, local area network security, and secure access to a file system, may all be implemented as applications. This can be less costly than having a special, secure subsystem for each task. Dial-in access can use standard modems, local area network security systems can use standard network adapters, and secure file systems can use standard disk controllers.

By reusing the resources of the protected processor for each of these tasks, the incremental cost of providing security for each of them can be reduced dramatically.

Conclusion

ABYSS is an architecture that enables the trusted execution of applications on protected processors, through its use of logical, physical, and procedural security. Software is separated from its Right-To-Execute, and strong cryptographic methods are used to manage both. Secure cryptographic channels can be used to move both software and Rights-To-Execute between protected processors.

Tokens are introduced as a new use-once authorization mechanism, which can be implemented with significantly less expense than the cryptographic facilities necessary to create secure channels. They are useful when authorizations are distributed physically, rather than electronically.

ABYSS does not require changes to current or planned software distribution methods. It is nearly transparent to the legitimate user. Software installation is automated, and there is no change in how software is executed. Rights-To-Execute software are not permanently fixed to the systems on which they are installed; they can be transferred between systems easily. Files containing the software may be stored on the user's hard disk, on network file servers, or on a mainframe, and may be backed up at will. Furthermore, Rights-To-Execute can themselves be backed up, to preclude loss in the event of a failure of the protected processor.

At the same time, all current and planned software distribution system are supported, and protected. Software authors and vendors may use the system to enforce the terms and conditions of their software sales technically. This opens up many new opportunities for marketing software.

Acknowledgements

The authors greatly appreciate the contributions of the ABYSS group, William C. Arnold, Thomas J. Nolan, Bill Strohm, and Steve H. Weingart, without whom these ideas would not have matured into working prototypes, of Akhileshwari Chandra, who suggested the privilege structure and helped with early incarnations of this architecture, and of Francis N. Parr, who helped simplify and refine several key features of the architecture.

References

- * The expression $E_K(M)$ represents a message M that has been encrypted under a key K . The cryptographic system used, and the mode in which it is used, may depend upon the situation in which they are used.
- ** The expression $\Omega(x)$ arises in complexity theory, and represents a quantity which (asymptotically) is at least as large as x . [Purd85]

- [Arno85] Mark G. Arnold, Mark D. Winkel, "Computer Systems to Inhibit Unauthorized Copying, Unauthorized Usage, and Automated Cracking of Protected Software," U.S. Patent No. 4,558,176, issued Dec. 10, 1985
- [Best79] Robert M. Best, "Microprocessor for Executing Enciphered Programs," U.S. Patent No. 4,168,396, issued Sept. 18, 1979
- [Best80] Robert M. Best, "Preventing Software Piracy With Crypto-Microprocessors," *Proc. IEEE Spring COMPCON 80*, Feb. 25-28, 1980, San Francisco, CA, pp. 466
- [Best81] Robert M. Best, "Crypto Microprocessor for Executing Enciphered Programs," U.S. Patent No. 4,278,837, issued July 14, 1981
- [Best84a] Robert M. Best, "Cryptographic Decoder for Computer Programs," U.S. Patent No. 4,433,207, issued Feb. 21, 1984
- [Best84b] Robert M. Best, "Crypto Microprocessor That Executes Enciphered Programs," U.S. Patent No. 4,465,901, issued Aug. 14, 1984
- [Blak85] G.R. Blakley, C. Meadows, G. B. Purdy, "Fingerprinting Long Forgiving Messages," in *Advances in Cryptology: Proc. Crypto 85*, Hugh C. Williams (ed.), pp. 180
- [Chau84] David Chaum, "Design Concepts for Tamper Responding Systems," in *Advances in Cryptology: Proc. Crypto 83*, David Chaum (ed.), Plenum Press 1984, pp. 387
- [Chur86] David Churbuck, "U.S. Software Firms Struggling to Stem Illegal Foreign Sales," in *PC Week*, Ziff-Davis Publishing Co., New York, pp. 213 (Nov. 11, 1986)
- [Cina86] Vincent J. Cina Jr., Steve R. White, Liam Comerford, "ABYSS: A Basic Yorktown Security System: PC Software Asset Protection Concepts," IBM Research Report Number RC 12401 (Dec. 18, 1986)
- [Denn83] Dorothy Elizabeth Robling Denning, *Cryptography and Data Security*, Addison-Wesley Publishing Co. (1983), pp. 192
- [Ever85] David Everett, "Padlock," *Computer Bulletin*, Ser. 3, No. 1, Pt. 1, pp.16 (March 1985)
- [Gold85] S. Goldwasser, S. Micali, C. Rackoff, "The Knowledge Complexity of Interactive Proof Systems," *Proc. 17th ACM Symp. on Theory of Computing*, pp. 291 (1985)
- [Gold86] Oded Goldreich, "Towards a Theory of Software Protection," *Proc. Crypto '86*, Santa Barbara, California, pp. 35-1 (1986)
- [Herz84] A. Herzberg, G. Karmi, "On Software Protection," *Proc. Fourth JCIT*, pp. 388 (April 1984)
- [Herz86] Amir Herzberg, Shlomit S. Pinter, "Public Protection of Software," in *Advances in Cryptology: Proc. Crypto 85*, Hugh C. Williams (ed.), pp. 158 (1986)
- [Juen83] R.R. Jeuneman, S.M. Matyas, C.H. Meyers, "Message Authentication with Manipulation Detection Codes," *Proc. 1983 Symp. on Security and Privacy*, pp. 33
- [Kent80] Stephen Thomas Kent, "Protecting Externally Supplied Software in Small Computers," Ph.D. Thesis, MIT Laboratory for Computer Science (Sept. 1980)
- [Pric86] W.L. Price, "Physical Security of Transaction Devices," NPL Technical Memo DITC 4/86, National Physical Laboratory (Jan. 1986)
- [Purd82] George B. Purdy, Gustavus J. Simmons, James A. Studier, "A Software Protection Scheme," *Proc. 1982 Symp. on Security and Privacy*, Oakland, California, pp. 99 (April 26-28, 1982)
- [Purd85] P.W. Purdom Jr., C.A. Brown, *The Analysis of Algorithms*, Holt, Rinehart and Winston (1985)
- [Sham79] A. Shamir, "How to Share a Secret," *Comm. ACM* 22, 11, pp. 612 (Nov. 1979)
- [Simm85] Gustavus J. Simmons, "How to (Selectively) Broadcast a Secret," *Proc. 1985 Symp. on Security and Privacy*, Oakland, California, pp. 108 (April 22-24, 1985)
- [Stro86] Bill Strohm, Liam Comerford, Steve R. White, "ABYSS Tokens," IBM Research Report Number RC 12402 (Dec. 18, 1986)
- [Voel86] John Voelker, Paul Wallich, "How disks are 'padlocked'," *IEEE Spectrum*, pp. 32 (June 1986)
- [Wein86] Steve H. Weingart, "Physical Security for the μ ABYSS System," *1987 IEEE Symp. on Security and Privacy*, these proceedings