



January 1997

A Secure Active Network Environment Architecture

D. Scott Alexander
University of Pennsylvania

William A. Arbaugh
University of Pennsylvania

Angelos D. Keromytis
University of Pennsylvania

Jonathan M. Smith
University of Pennsylvania, jms@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith, "A Secure Active Network Environment Architecture", . January 1997.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-97-17.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/114
For more information, please contact repository@pobox.upenn.edu.

A Secure Active Network Environment Architecture

Abstract

Active Networks are a network infrastructure which is programmable on a per-user or even per-packet basis. Increasing the flexibility of such network infrastructures invites new security risks. Coping with these security risks represents the most fundamental contribution of Active Network research. The security concerns can be divided into those which affect the network as a whole and those which affect individual elements. It is clear that the element problems must be solved first, as the integrity of network-level solutions will be based on trust of the network elements.

In this paper, we describe the architecture and implementation of a Secure Active Network Environment (SANE¹), which we believe provides a basis for implementing secure network-level solutions. We guarantee that a node begins operation in a trusted state with the AEGIS secure bootstrap architecture. We guarantee that the system remains in a trusted state by applying dynamic integrity checks in the network element's run time system, a novel naming system, and applying node-node authentication when needed.

The SANE implementation is for x86 architectures, currently those running one of several varieties of UNIX.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-97-17.

A Secure Active Network Environment Architecture

D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis and Jonathan M. Smith

“Доверяй, но Проверяй”

“Trust, but Verify”

Abstract—

Active Networks are a network infrastructure which is programmable on a per-user or even per-packet basis. Increasing the flexibility of such network infrastructures invites new security risks. Coping with these security risks represents the most fundamental contribution of Active Network research. The security concerns can be divided into those which affect the network as a whole and those which affect individual elements. It is clear that the element problems must be solved first, as the integrity of network-level solutions will be based on trust of the network elements.

In this paper, we describe the architecture and implementation of a Secure Active Network Environment (SANE¹), which we believe provides a basis for implementing secure network-level solutions. We guarantee that a node begins operation in a trusted state with the AEGIS secure bootstrap architecture. We guarantee that the system remains in a trusted state by applying dynamic integrity checks in the network element’s run time system, a novel naming system, and applying node-node authentication when needed.

The SANE implementation is for x86 architectures, currently those running one of several varieties of UNIX.

I. INTRODUCTION

A variety of proposals for programmable network infrastructures are currently extant, such as open signaling [1] and Active Networks [2]. These proposals share the goal of improving network flexibility and functionality through introduction of an accessible programming abstraction, which may be available on a per-user or even a per-packet basis. In the SwitchWare project [3], U. Penn and Bellcore are collaborating on research into the architecture of Active Network elements.

The goal of programmable network architectures is to provide an acceleration of network service creation. Protocols provide a set of rules by which compliant systems can participate in communications. To build a global virtual infrastructure such as the IP [4] Internet, a “minimal” interoperability requirement was set, namely a packet format and a common addressing scheme. The IP hourglass is shown in Figure 1. Service enhancements, such as the TCP reliable stream protocol [5], occur at the end-points of the virtual infrastructure. Since all IP-compliant network infrastructures must support the IP protocol, change of the infrastructure itself is slow and highly constrained. As the Internet has become commercialized, the standardization process has slowed considerably; yet at the same time there is increasing demand for enhanced services.

Active Networks follows the approach first proposed in the “Protocol Boosters” project [6], of enabling on-the-fly modification of network functionality, for example to adapt to changes in link conditions. Protocol Boosting is a design methodology, but Active Networks provides an infrastructure general enough

Scott Alexander, William Arbaugh, and Angelos Keromytis are each working toward a Ph.D. in computer and information science at the University of Pennsylvania.

Jonathan M. Smith is an associate professor at the University of Pennsylvania.

Old Russian saying.

¹The pun for InSecure Active Network Environment is obvious (and painful).




Fig. 1. Hourglass model of internetworking.

to support any network reprogramming. This is done by raising the level of abstraction of the interoperability layer from a packet format to a programming environment accessible to programmers. Not surprisingly, there are applications for a programmable network infrastructure.

A. Applications of Active Networks

There are many applications of programmable network infrastructures, some of which we can imagine today, and some which will only appear as the limitations and implications of the technology are discovered. We give here three simple examples of uses which would (1) enhance flexibility, (2) improve performance over today’s networks, and (3) improve manageability relative to today’s networks.

- Since the virtual infrastructure provided by IP provides both forwarding and routing, these services are not subject to user control. This is a problem if one desires value-added services such as non-co-routed paths (to enhance throughput via striping or reliability in the face of link failure). With a programmable network infrastructure, a larger portion of the network connectivity might be usefully employed for enhanced services. The most important factor is that this can be done under user control.

- Congestion remains a major problem for our information infrastructures, especially as the number of connected nodes has exploded without corresponding increases in the core capacity. Using congestion pricing and e-cash like schemes, we can use price as a priority mechanism to spread load and let all available capacity serve as a queuing sink. Per-packet programs can let packets make their own routing decisions based on available information and resources, much as an automobile driver reroutes based on deadlines and traffic reports on the car radio. In this way, distributed intelligence, low-latency decision-making (*i.e.*, close to the congestion point) and economic algorithms (which are very scalable) can be employed in computer networks.

- Loadable diagnostic functionality is very powerful in network

management for a number of reasons. First, it is of course flexible. Second, it is only loaded when needed and hence need not be resident on any “fast path”. Third, it allows, when inserted in a multiplicity of network elements, a level of distributed monitoring (*e.g.*, for intrusion detection and other tasks) that has heretofore proven difficult if not impossible. For example, the approach used in Paxson’s thesis [7] of randomized traceroute probes illustrates the difficulty of obtaining routing statistics.

B. Threats

Threats to network infrastructure are intimately tied to the model used for sharing the infrastructure. For example, with the unreliable transport model provided by the Internet, any security policies are enforced end-to-end. To the infrastructure, packets are anonymous; only the destination address is used, in concert with a routing algorithm, to select an entry from the forwarding table.

IP packets are anonymous to the routers, and they, at least before extensions such as MBONE [8] and RSVP [9], are allocated service on a FIFO basis. IPSEC [10] provides authentication services, but it remains unclear how support for Quality of Service (such as RSVP) will be integrated with authentication services. As it stands, the Internet infrastructure is vulnerable to a variety of denial of service attacks as a consequence of minimal resource accountability, as well as a variety of other attacks such as traffic analysis. We note that since the resource model in the routers is so simple, sophisticated threats are posed by attacks on services implemented at the endpoints, *e.g.*, the notorious “Syn-Ack” attack [11] on TCP/IP and the “Ping of Death” [12].

Active Networks, being more flexible, considerably expand the threat possibilities. The security threats faced by such elements are considerable. For example, when a packet containing code to execute arrives, the system typically must:

- Identify the sending network element
- Identify the sending user
- Authorize access to appropriate resources based on these identifications
- Allow execution based on the authorizations and security policy

In networking terminology, the first three steps comprise a form of admission control, while the final step is a form of policing. A second view is that of static versus dynamic checking. Security violations occur when a policy is violated, *e.g.*, reading a private packet, or exceeding some specified resource usage.

C. A high-level view of a SANE architecture

Systems are organized as layers to limit complexity. A common layering principle is the use of levels of abstraction to mark layer boundaries. A computer system is organized in a series of abstraction levels, each of which defines a “virtual machine” upon which higher levels of abstraction are constructed. Each of the virtual machines presupposes that it is operating in an environment where the abstractions of underlying layers can be treated as axiomatic. When these suppositions are true, the system is said to possess *integrity*. Without integrity, no system can be made secure.

Thus, any system is only as secure as the foundation upon which it is built. For example, a number of attempts were made in the 1960s and 1970s to produce secure computing systems using a secure operating system environment as a basis [13]. An essential presumption of the security arguments for these designs was that the system layers underpinning the operating system, whether hardware, firmware, or both, were trusted. We find it surprising, given the great attention paid to operating system security [14] [15] that so little attention has been paid to the underpinnings required for secure operation, *e.g.*, a secure bootstrapping phase for these operating systems.

In a computer system, the integrity of lower layers is typically treated as axiomatic by higher layers. Under the presumption that the hardware comprising the machine (the lowest layer) is valid, the integrity of a layer can be guaranteed *if and only if*: (1) the integrity of the lower layers is checked, and (2) transitions to higher layers occur only after integrity checks on them are complete. The resulting integrity “chain” inductively guarantees system integrity. We call this the Chaining Layered Integrity Checks (CLIC) model.

The overall approach to security taken in the SwitchWare project is to provide carefully circumscribed functionality to network programmers, by means of a programming language which allows us to limit functionality and run in a controlled environment. We have implemented a prototype of such a network element, and applied it to the problem of constructing an extended LAN (bridging).

II. SANE ISSUES AND ARCHITECTURE

In this section we discuss the issues which arise from the threat model we presume. After a discussion of these issues, we further discuss integrity and trust relationships at various levels in the system. Finally, we give a high-level architecture which addresses the division of integrity checking and enforcement into static and dynamic portions.

A. Separation of Concerns

We make a somewhat artificial, albeit useful, division of our concerns into *static* and *dynamic*. Static concerns are those which can be checked once, or infrequently, as in the case of an active network bootstrapping from a cold start into an operational state. Dynamic concerns are those which must be continuously addressed to maintain the operational state of the system.

There are several major advantages to this division that can be used in a system design. First, as static checks are done once, or very few times, they can be very expensive if this pays off in a significant increase in security. Second, dynamic checks can be made faster if it is known that the static checks have been performed in advance. Finally, these divisions usually closely follow the division of a system into layers of abstraction. If the proper trust and integrity relationships are preserved, the operation of the entire system can be trusted.

B. Integrity and Trust

Integrity is a way of saying that a system is what we expect it is; that is, it is unmodified. Trust is a more complex relationship, as something can be unmodified, but not trusted, while if

a system is trusted, it must remain unmodified for the trust relationship to hold.

Integrity and Trust relationships in an active network setting are of several types. In a layered architecture, each layer in a system trusts the layer below it.

For an active network node, a trusted node architecture can be constructed by making the lowest layers of the system trusted, and then ensuring that higher layers depend on the integrity of these lower layers.

There is, however, a significant difference when the actions of the node are programmable and the programs come from hosts or other active network elements. In this case, we must construct a web of trust between participating elements. Further, trust is not enough: a downloaded program from a trusted node may be flawed and may damage the receiving node. *Dynamic integrity* checks ensure that the node remains a participating element of the active network in spite of such threats.

It is clear then that any architecture for system security in an Active Network must use a combination of static checks and dynamic checks to remain secure.

C. Architecture

The basic layered structure of the Secure Active Network Environment (SANE) is shown in Figure 2. Here, we will explain the overall organization of the architecture and its principal goals. The remainder of this article expands on the components of the architecture.

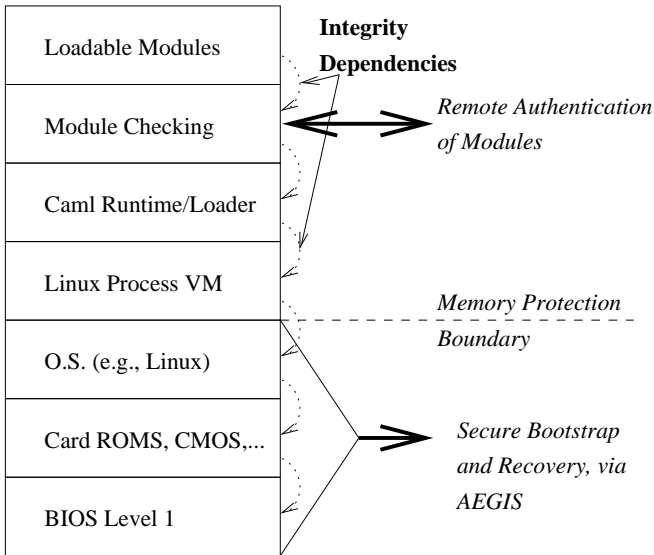


Fig. 2. SANE Architecture

The lower layers of the architecture ensure that the system *starts* in an expected state. The design utilizes a secure bootstrap architecture, called AEGIS, to reach the stage where dynamic integrity checks can be applied on a per-user or per-packet basis. AEGIS assumes the integrity of the system level-1 BIOS, and little else². It then repeatedly, until the operational active network element is operating, checks the integrity of the succeeding layer in the bootstrap before passing control to it. In-

²For recovery, a network-accessible trusted source is required.

tegrity is checked with a digital signature. This process results in the *expected* operational system starting execution; it makes no guarantees that that system operates correctly. Eventually, we hope to address at least a fraction of operational correctness issues with the application of formal methods.

When the active network element system is operational, it maintains security in several ways. First, it performs remote authentication when required for node-to-node authentication. Second, it provides a restricted execution environment for the evaluation of switchlets (the programs received from the network). Finally, it uses a novel naming scheme we have developed to partition the node's services name space between users. The authentication and integrity checks performed before a language system begins operating on it, such as checking a digital signature, are static. This is in contrast to dynamic checks performed (*e.g.*, by trying to type-check the packet's code or constrain its execution). These latter checks are performed frequently and thus must be performed efficiently; they guarantee that the network element remains secure, and remains operating.

D. Public Key Infrastructure

A very important element of our proposed architecture is the public key infrastructure. It is assumed that every user (or group of users) and every active element own a public/private key pair, and that these keys (and certificates) are used to authenticate and authorize actions of those entities³. It is also desirable that the infrastructure allows selective authorization delegation, so that flexible access and resource control policies can be built. Finally, depending on the underlying network fabric, the preferable method to revoke a certificate is by expiration; this minimizes network traffic when authorization checks are performed. In our implementation we intend to use SPKI [16] and Policy-Maker [17].

III. AEGIS ARCHITECTURE

An ideal CLIC would work with each level verifying the next as represented by the recurrence shown in equation 1.

$$I_0 = True, \quad (1)$$

$$I_{i+1} = \left\{ I_i \wedge V_i(L_{i+1}) \quad \text{for } 0 < i \leq n, \right.$$

where I_i represents a boolean value of the integrity of level i , n represents the number of levels in the bootstrap, and \wedge is the boolean *and* operation. V_i is the verification function associated with the i^{th} level. V_i takes as its only argument the level to verify, and it returns a boolean value as a result. Unfortunately, implementing the recurrence in equation 1 is difficult, if not impossible, in current computer systems.

Without a secure bootstrap the evaluator running on the network element cannot be trusted since it is invoked by an untrusted process.

A. AEGIS Overview

AEGIS modifies the boot process shown in Figure 3 so that all executable code, except for a very small section of trusted code, is verified prior to execution by using a digital signature. This is

³Key owners will be referred to as principals.

accomplished through modifications and additions to the BIOS (Basic Input Output System). The BIOS contains the verification code, and public key certificate(s). In essence, the trusted software serves as the root of an authentication chain that extends to the evaluator and potentially beyond to “active” packets. In the AEGIS boot process, either the active network element is started, or a recovery process is entered to repair any integrity failure detected. Once the repair is completed, the system is restarted to ensure that the system boots. This entire process occurs without user intervention.

In addition to ensuring that the system boots in a secure manner, AEGIS can also be used to maintain the hardware and software configuration of a machine. Since AEGIS maintains a copy of the signature for each expansion card⁴, any additional expansion cards will fail the integrity test. Similarly, a new evaluator cannot be started since the boot block would change, and the new boot block would fail the integrity test.

B. AEGIS Boot Process

Every computer with the IBM PC architecture follows approximately the same boot process. We have divided this process into four levels of abstraction (see Figure 3), which correspond to phases of the bootstrap operation. The first phase is the Power on Self Test or POST [18]. POST is invoked in one of four ways:

1. Applying power to the computer automatically invokes POST causing the processor to jump to the entry point indicated by the processor reset vector.
2. Hardware reset also causes the processor to jump to the entry point indicated by the processor reset vector.
3. Warm boot (*ctrl-alt-del* under DOS) invokes POST without testing or initializing the upper 64K of system memory.
4. Software programs, if permitted by the operating system, can jump to the processor reset vector.

In each of the cases above, a sequence of tests are conducted. All of these tests, except for the initial processor self test, are under the control of the system BIOS.

Once the BIOS has performed all of its power on tests, it begins searching for expansion card ROMs which are identified in memory by a specific signature. Once a valid ROM signature is found by the BIOS, control is immediately passed to it. When the ROM completes its execution, control is returned to the BIOS.

The final step of the POST process calls the BIOS operating system bootstrap interrupt. The bootstrap code first finds a bootable disk by searching the disk search order defined in the CMOS. Once it finds a bootable disk, it loads the primary boot block into memory and passes control to it. The code contained in the boot block proceeds to load the operating system, or a secondary boot block depending on the operating system [19] [20] or boot loader [21].

Ideally, the boot process would proceed in a series of levels with each level passing control to the next until the operating system kernel is running as modeled by Equation 1. Unfortunately, the IBM architecture uses a “star like” model which is shown in Figure 3.

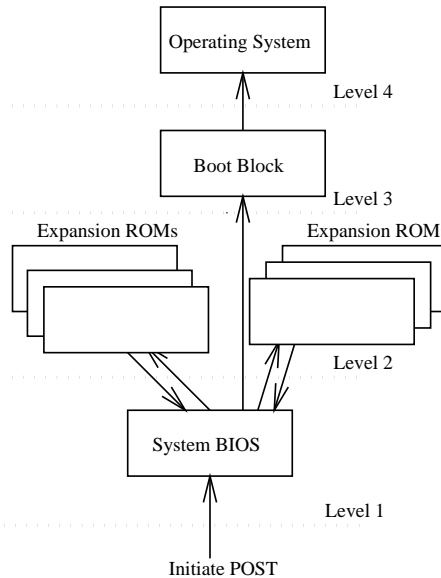


Fig. 3. IBM PC boot process

B.1 A Layered Boot Process

We have divided the boot process into several levels to simplify and organize the AEGIS BIOS modifications, as shown in Figure 4. Each increasing level adds functionality to the system, providing correspondingly higher levels of abstraction. The lowest level is Level 0. Level 0 contains the small section of *trusted* software, digital signatures, public key certificates, and recovery code. The integrity of this level is assumed to be valid. We do, however, perform an initial checksum test to identify PROM failures. The first level contains the remainder of the usual BIOS code, and the CMOS. The second level contains all of the expansion cards and their associated ROMs, if any. The third level contains the operating system boot block(s). These are resident on the bootable device and are responsible for loading the operating system kernel. The fourth level contains the operating system, and the fifth and final level contains user level programs and any network hosts.

The transition between levels in a traditional boot process is accomplished with a jump or a call instruction without any attempt at verifying the integrity of the next level. AEGIS, on the other hand, uses public key cryptography and cryptographic hashes to protect the transition from each lower level to the next higher one, and its recovery process ensures the integrity of the next level in the event of failures. The pseudo code for the action taken at each level, L , before transition to level $L + 1$ is shown in Figure 5. The function *IntegrityValid* first finds the component certificate for Level l . Ideally this will be stored in the component itself, but initially it will be stored in a table contained in Level 0. Once the certificate, c , is found. *VerifyCertChain* then verifies that the certificate(s) form a “chain” of trust from the component certificate to the root Certificate Authority Public Key. If they do not, then both *VerifyCertChain* and *IntegrityValid* return FALSE and a recovery procedure is entered. If *VerifyCertChain* returns TRUE, then the signature contained in the certificate is verified using the public key contained in the certificate.

⁴Ideally, the signature would be embedded in the firmware of the ROM.

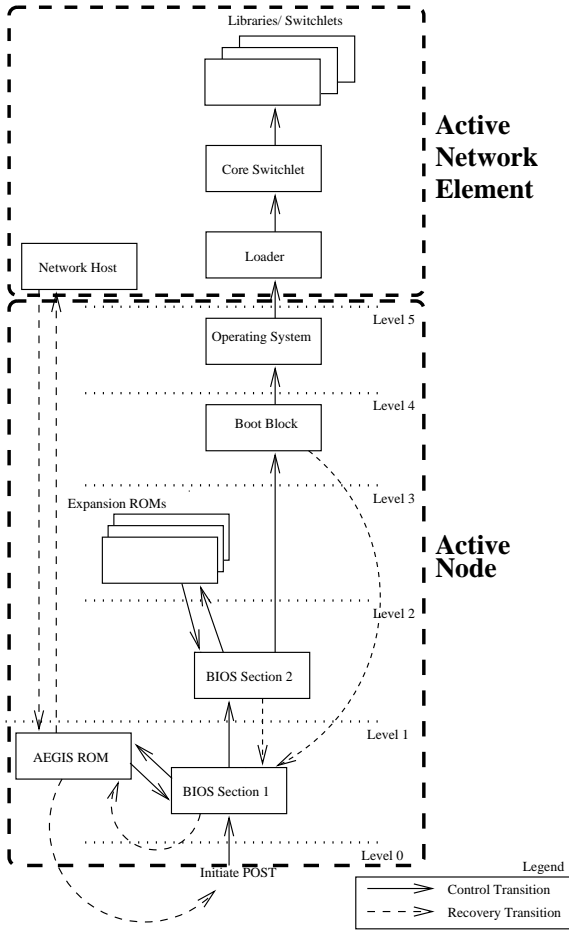


Fig. 4. AEGIS boot control flow

B.2 AEGIS BIOS Modifications

AEGIS modifies the boot process shown in Figure 3 by dividing the BIOS into two logical sections. The first section contains the bare essentials needed for integrity verification and recovery. It comprises the “trusted software”. The second section contains the remainder of the BIOS and the CMOS.

The first section executes and performs the standard checksum calculation over its address space to protect against ROM failures. Following successful completion of the checksum, the *IntegrityValid* function is called. If the function returns true, then control is passed to the second section, *i.e.*, Level 1.

The second section proceeds normally with one change. Prior to executing an expansion ROM, the function *IntegrityValid* is called. If the function returns true, then control is passed to the expansion ROM. Once the verification of each expansion ROM is complete (Level 2), the BIOS passes control to the operating system bootstrap code. The bootstrap code was previously verified as part of the second section of the BIOS or expansion ROM, and thus no further verification is required. The bootstrap code finds the bootable device and verifies the boot block.

Assuming that the boot block is verified successfully, control is passed to it (Level 3). If a secondary boot block is required, then it is verified by the primary block before passing control to it. Finally, the kernel is verified by the last boot block in the chain before passing control to it (Level 4). This results in the

```

Boolean IntegrityValid(Level l) {
    Certificate c = LookupCert(l);

    if (VerifyCertChain(c))
        return DSAVerify(SHA1(l), c);
    else return FALSE;
}

if (IntegrityValid(L+1)) {
    GOTO(L+1);
} else {
    GOTO(Recovery);
}

```

Fig. 5. Layer Transition Pseudo code

CLIC model shown in Equation 2 and the layering model shown in Figure 4. Any integrity failures identified in the above process are recovered through a trusted repository.

$$\begin{aligned}
 I_0 &= \text{True}, \\
 I_{i+1} &= \begin{cases} I_i \wedge V_i(L_{i+1}) & \text{for } i = 0, 3, 4, \\ I_i \wedge \sum_{l=1}^n V_i(L_{i+1}^l) & \text{for } i = 1, \\ I_i \wedge V_{i-1}(L_{i+1}) & \text{for } i = 2. \end{cases} \quad (2)
 \end{aligned}$$

C. Recovery Process

The trusted repository can either be an expansion ROM board that contains verified copies of the required software, or it can be a another active node. If the repository is a ROM board, then simple memory copies can repair or shadow failures. If the repository is a network host, then a protocol with strong authentication is required. We describe this protocol in Section III-C.1.

In the case of a network host, the detection of an integrity failure causes the system to boot into a recovery kernel contained on the network card ROM. The recovery kernel contacts a “trusted” host through the secure protocol described in this paper to recover a signed copy of the failed component. The failed component is then shadowed or repaired, and the system is restarted (warm boot).

C.1 Recovery Protocol

The protocol we use throughout this paper and in our architecture is based on the Station to Station protocol [22]. The basis of the protocol is the Diffie-Hellman exchange [23] for key establishment, and public key signatures for authentication (to avoid man in the middle attacks). In our architecture we use DSA [24] digital signature standard, but other (RSA [25] etc.) algorithms can be used.

Briefly, this protocol allows each participant to establish the identity of the other, discover the operations that the peer is authorized to perform, and allows the two parties to establish a shared secret to be used for a variety of purposes including the authentication and encryption of future traffic. This is accomplished by having each party send the other both an authentication certificate and an authorization certificate and using Diffie-Hellman key exchange to establish the shared secret. The protocol is carried out with a total of three messages transmitted.

A node that has detected an integrity failure can establish this trust relationship with a repository. It can then request a new version of the failed component. The repository will send the new component using DSA to guarantee that this is the correct component and using the shared secret to ensure that the component is not tampered with while transiting the network.

Authentication Exchange Protocol

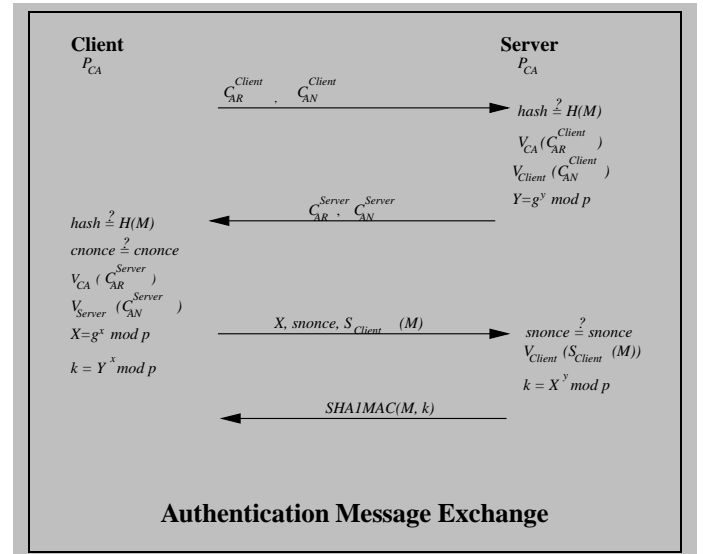
In more detail, here's a description of the authentication exchange protocol. The initiator sends a message to the other party containing its authorization and authentication certificates. The responder receives the message and verifies the initiator's signature on the authentication certificate and that the hash contained in the authentication certificate matches that of the message, M . The authorization certificate is also verified^a.

If all are valid and the timestamp on the authentication certificate is within bounds, then the responder sends to the initiator a message containing its authorization and authentication certificates. The responder's authentication certificate may include the optional DH parameters, g and p , and Y , where $Y = g^y \text{ mod } p$. If the DH parameters are not included in the certificate, then default values for g and p are used. Currently, we are using the same default values as those used in SKIP [26]. The responder's nonce, $snonce$, and the initiator's nonce, $nonce$, are also included in the message. The initiator receives this message and verifies the signatures on the authentication and authorization certificates, that the hash in the responder's authentication certificate matches the message hash, and that $nonce$ matches that sent in the first message.

If all are valid and the timestamp value of the authentication certificate is within bounds and $nonce$ matches that sent in the first message, then the initiator sends a signed message to the responder containing its DH parameter X where $X = g^x \text{ mod } p$, and the responder's nonce $snonce$. The responder receives the message and verifies the signature and that $snonce$ matches that sent in its previous message. If both are valid, then the responder can generate the shared secret, k , using DH. The initiator similarly generates k . Now, k , can be used to authenticate messages between the initiator and the responder until such time as both agree to change k . The figure depicts the entire exchange between the initiator and the responder.

The use of the authentication certificate assists in ensuring that the protocol is "Fail Stop" [27] through the use of nonces and a short validity period for the certificate. The use of $snonce$ also permits the responder to reuse Y over a limited period. This reduces the computational overhead on the responder (who will typically be an active node) during high activity periods. The potential for a Syn-Ack like denial of service attack is mitigated in the same manner by the authentication certificate.

^aThis implies that there is a certificate chain from the responder's key to the initiator's certificate. Trusted third parties can be part (or the beginning) of such a chain.



IV. BOOTSTRAPPING A SANE NETWORK

Once the node has been brought up in a secure manner, it attempts to establish trust relations with its direct peers. The same protocol that was described in section III-C.1 is used to exchange certificates and establish a shared secret key with each of the peer active nodes. The certificates exchanged at this stage are used to verify the neighbors, establish administrative domains (and their boundaries) and the trust relations inside and between those domains. The secret key and the trust relations will then be used to:

- Minimize path setup costs, as we'll describe in the section VI-A.
- Allow mobile-agent [28] [29] type of applications, where per-hop authentication (and possibly encryption) may be necessary. An API will be defined that lets a programmer make use of these services.
- Secure message exchange between peer active nodes, such as for routing messages or network management.
- Establish authenticated packet forwarding channels.
- Deter link traffic analysis; the active node administrator will then be able to allocate a percentage of the available bandwidth as an encrypted, always-busy, channel. An eavesdropper on the link will then be unable to determine which messages were forwarded to the peer node. Again, an API will be defined that programmers can take advantage of.

V. THE ACTIVE NETWORK INFRASTRUCTURE

With the operating system verified and booted, the next step is to make the node part of the active network. This is accomplished by loading two final layers. Given our definition of active networks, not surprisingly, the lower layer of our network infrastructure is a loader which can load switchlets, our active programs. On top of the loader is a Core Switchlet which provides essential services. Finally, a non-privileged layer consisting of a set of library routines which provide common services is added. This layering, together with the applications or switchlets, is illustrated in Figure 4.

The lower two layers provide the basis of the dynamic security model in the network infrastructure. They do this by using a

strongly-typed language which supports garbage collection and module thinning. Using these techniques, we move from static to dynamic enforcement of our security mechanisms.

A. Why Does the Language Matter?

The programming language defines what operations the programmer can perform. By careful choice of language, we can limit some of the undesirable actions that a programmer might unintentionally or maliciously perform. Thus, through the choice of language, we can prevent certain classes of security violations.

The first property that we desire from the language is strong typing. In a strongly typed language, the only way to convert data from one type to another is through a well-defined conversion routine. Thus, one can typically transform an integer into a floating point value, but cannot perform conversions to or from a pointer type. In a weakly typed language like C, it is this ability to freely convert types which leads to the need for heavier security mechanisms including separation of address spaces between processes.

The second property that we desire is garbage collection. If the programmer is able to manage storage directly, two problems can result. The first is failure to free storage which can lead to loss of performance throughout the system. The second, more dangerous problem, occurs when storage is returned to the allocator and then referenced later. If the storage has been reassigned to another user, it is possible to discover another user's information. Worse yet, if the address is no longer valid, a fault results which must be handled to avoid crashing the entire system.

The third property that we desire is module thinning. By modules, we mean a set of functions and values which have been combined into a package by the programmer. Module thinning is a technique which allows us to pick and choose which functions and values from a module are available to a switchlet which we load. For example, in the Thread module that we use, there is a function which allows one to kill any program on the system by specifying its process ID. This is inappropriate for switchlets, so we do not make this available except to the loader and the Core Switchlet.

The final property which we require is the ability to dynamically load programs. Clearly, if we intend to run programs that arrive over the net, we must have a way to link those programs into the running system and evaluate them. Dynamic loading gives us this ability.

The Caml programming language [30] provides these features. Caml additionally provides us with a threads interface and static type checking. The former allows a natural programming style and precludes the need to implement a scheduler. The latter pushes many of the costs associated with the type system to compile time. Thus, checks that other systems perform repeatedly at runtime, we perform once at compile time.

B. The Loader

The loader forms the basis of the dynamic security for our network infrastructure. Once it has been securely started by the AEGIS bootstrap, the loader provides a minimal set of services necessary to find the Core Switchlet and start it running. It also

provides policy and mechanism for making changes to the Core Switchlet, if that is desirable.

The loader is also responsible for providing the mechanism by which modules are loaded. Currently, the mechanisms provided are loading from disk or loading from a string. The Core Switchlet governs the policy by which this mechanism may be used and may provide interfaces to the mechanism.

C. The Core Switchlet

The Core Switchlet is the privileged portion of the system visible to the user. Through the use of module thinning, it determines which functions and values are visible to which users. The services that it provides are broken into five modules.

The first module is `Safestd`. This module provides the functions that one would expect to find in any programming language including addition and multiplication as well as more complex abstractions like lists, arrays, and queues. Many functions including the I/O functions have been thinned from this module to make it safe.

The next module is `Safeunix`. This module has been very heavily thinned; it gives access to Unix error information, some time related functions, and some types that are needed for the networking interface that we provide. The rest of the access to Unix functions has been thinned away.

In order to allow the user to supply error or status messages, we have a `Log` module. The user supplies a string which will be saved to a system log. What and where this system log is, is not defined. For convenience while debugging, we currently write the messages to a disk file, but for security purposes, we intend to extend this module to limit the amount and frequency of messages produced by any given thread.

Access to the network is provided by the `Unixnet` module. This allows switchlets to access network interfaces for either sending or receiving frames. Currently, only one switchlet is allowed to have access to a given interface. In the near future, we intend to modify this module to receive and demultiplex the data. Access to the data will then be available to any switchlet, assuming said switchlet can prove that it has the authority to access the data as described in section VI-A.

The last of the five modules is `Safethread`. As mentioned, this provides a threads package which helps in the structuring of the system. Each switchlet runs in a thread and is capable of creating additional threads. When a switchlet is first started, it is given an identifier inside of an opaque type. (An opaque type is one which has no conversion functions to or from any other type. Thus, the identifier cannot be forged.) In order to use additional resources including creating additional threads, the switchlet must provide its identifier which allows the system to check the resources currently consumed and allow or deny the request for additional usage.

D. The Library

The library is a set of functions which provide useful routines which do not require privilege to run. The proper set of functions for the library is a continuing area of research. Some of the things that are in the library for the experiments that we have performed include utility functions and implementations of IP and UDP [31].

E. The Active Bridge: An Active Networking Application

To demonstrate the utility of this infrastructure, we have implemented an Active Bridge. This bridge is built from several switchlets which build up layers of functionality. In particular, by loading just the lowest layer, we can demonstrate a buffered repeater. The next switchlet adds a learning algorithm. Finally, the highest layer of the bridge adds spanning tree functionality to give us a nearly compliant [32] bridge.

VI. DYNAMIC SECURITY CHECKS

Once the active node is operating, we rely upon dynamic security checks and measures to ensure that the access and resource use policies defined by the administrator are followed. Furthermore, the node needs to provide certain guarantees in regards to service access⁵; some of these guarantees are provided by the underlying operating system and programming language. However, some of our guarantees must be built through additional mechanisms provided by our system.

A. Access Control

One of the basic goals of active networks is allowing users to install their own protocols on network elements, in the form of dynamically loaded modules. Since these modules may have access to critical resources, it is imperative that this access be controlled. Furthermore, in some cases it is necessary to authenticate packets belonging to some particular packet sequence, if they need to be handled in some “privileged” manner (*e.g.*, going through a firewall or delivery to some service). In the next two sections, we extend the mechanism described in section III-C.1 to provide authentication and authorization mechanisms.

These trust relations will be established along a path of active nodes in most cases. Two possible methods of path establishment are:

- Via direct negotiation with each node, possibly in parallel. The implication here is that the initiator can both identify and communicate directly with these nodes, instead of having to discover the path.
- In a “telescopic” manner, in which a scout packet would identify the next node at each step and initiate the negotiation. In this model, each negotiation has to finish before the next one begins (in order to establish a communications path from the current node back to the initiator).

A.1 Principal Authentication/Authorization

On an active node, when a principal requests an action (such as use a resource) that is privileged according to local policy, he has to provide credentials that authorize him to perform said action. The protocol that would implement the negotiation is the modified version of the STS protocol, as described in section III-C.1.

Once the node and the principal have established a security association, they can use it to authenticate (and possibly encrypt) all or some of the messages between them. The node retains all the credentials associated with this exchange, so it can determine whether future attempted actions of the principal are acceptable.

⁵Essentially, user isolation.

Figure 6 shows the packet format once the security association is established. The authenticator will be included in the packet along with an SPI⁶ and a replay detection counter, similar to the IPsec Authentication Header [33].




Fig. 6. Authenticator Header

However, doing this negotiation with every node along a path to a remote end node is bound to prove costly in two ways:

1. time (both real and CPU cycles spent on the cryptographic operations)
2. more importantly, packet overhead; for every node in the path, there would have to be a different authenticator (since the shared key is different between the principal and each node)

The impact of these problems and their solutions depends on the environment in which the active node is operating. Based on the types of attacks which must be protected against, we describe a series of measures which may be taken.

The first step is a simple optimization; once the described negotiation has taken place, the principal can then use the shared key to distribute another secret key to all the nodes in the path. By using this common key it is possible to have only one authenticator in the packet, which would be verifiable by all the nodes in the path.

There are two potential problems with this approach. There is still significant computational and path establishment overhead. If connections tend to be reasonably long running, this cost will be dwarfed.

A worse problem occurs because a (malicious) node in the path can perform actions as if it were the principal, on another node, since the key is shared between all the nodes. There are a few workarounds to this problem. In some environments, it may be adequate to accept the problem and to establish paths only through trusted nodes. This is likely to be impossible in other environments.

A second workaround is to distinguish between packet authentication and privileged operations. Authentication can be done using the common key, while privileged operations have to make use of the key known only to the particular node and the principal. This means that control operations will be safe, but that “data” can be forged or modified by a malicious or malfunctioning node. Finally, to avoid delivery of “bad” data to the remote endpoint, the packet would then have a second authenticator in it, which would be only verifiable by the two endpoints (and hence be unforgeable⁷ by intermediate nodes). If it is im-

⁶The SPI is a value used along with the principal and/or node identifiers to indicate the particular security association.

⁷Keeping in mind the assumptions made about the strength of the key and the algorithm used, of course.

portant to not deliver corrupted packets to modules running in intermediate nodes, there is a certain probabilistic scheme that can be used to detect tampering, described in Appendix A.

A last optimization is possible, by taking into consideration the results of section IV. If active nodes in the same administrative domain have a common set of policies regarding access control and resource utilization, it may be sufficient to go through the negotiation protocol once for each such domain (when entering it), and then having the credentials forwarded as necessary, as shown in Figure 7. This reduces the computational effort and the packet overhead necessary to authenticate/authorize the principal and subsequent packets.

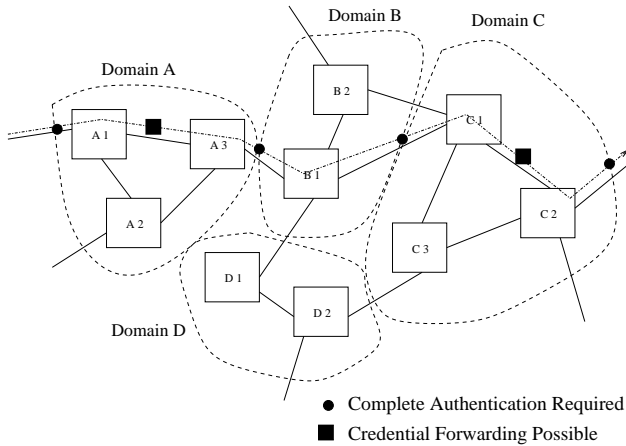


Fig. 7. Administrative Clouds and Path Setup

The above optimizations can be applied when the policy specified by the owner of the packet flow and the individual node policies allow them. Access policy to a module (once it has been successfully loaded) is specified by the module itself. The active node will then enforce this policy.

A.2 Single Packet Authentication

For certain classes of applications, the initiating principal may not know exactly which nodes an active packet will visit (e.g., mobile agent type of applications). This means that security association negotiation, as described in the previous section, may not be feasible. However, these programs may need to perform privileged operations on the active nodes, which means that some form of security guarantees has to be provided. There are a few approaches that can be taken:

- If the administrative domains through which the switchlet will travel are known *a priori*, the initiating principal may establish security associations with nodes in those domains. The established trust relations described in section IV can then be used to forward the credentials inside those domains.
- If the switchlet does not need to perform any privileged operations but requires some security guarantees of its own, it can make use of the existing peer to peer trust relations to do per-hop authentication and/or encryption. For example, if a switchlet requires that each node in its path belong to a list of nodes that it trusts, and since it must trust its creating node, at any hop, it is on a trusted node and can request that that node use its security relations to forward it to another node that is on its list of trusted nodes.

- The switchlet can carry all necessary authorizations the initiating principals believes it may need. These authorizations would be in the form of public key certificates, and the agent needs to be authenticated through a digital signature. While this approach is quite simple, it has two primary drawbacks. It wastes packet space, since all the certificates need to be carried even if they are not used. Further, it is hard to avoid switchlet-replay, unless we assume either network wide (roughly) synchronized clocks or persistent state on the active nodes⁸. Providing these allows a series of potential denial of service attacks.
- When the switchlet needs to perform some privileged operation and needs credentials, it can notify the initiating principal who can then initiate a negotiation to establish a security association. Credentials can then be carried along while inside the same administrative domain. The assumption here is that the switchlet is able to send the notification message back to the principal, which depends on the both the underlying network infrastructure and the node policies.

VII. DYNAMIC RESOURCE NAMING

Conceptually, loaded modules can be considered as the interfaces to user defined resources. Such resources will generally be shared between different sessions of the same principal, or even between different principals. These principals will need to identify (name) the particular resource they want to use.

The “naive” way of naming (using some user-defined value) would not work well, because names need to be unique across the active network. If users arbitrarily assign names to their resources, it is conceivable that there will be accidental naming collisions; worse yet, forging names is possible, allowing for resource-in-the-middle attacks. Alternatively, some centralized authority could assign names per request, making sure these remain unique; this solution is unattractive because it does not scale well as the number of names required increases.

We present a decentralized way of naming dynamic resources that does not allow name collisions, accidental or malicious. We are one assumption: in order to load a module on the active element, the principal must pass some type of authorization check. Furthermore, this authorization is fine grained; each principal is distinguishable for our purposes⁹. We believe that this assumption is reasonable, since we expect that an active element owner will probably want to limit the resources that any principal will potentially consume. (Moreover, we expect that the owner will want to give different access and resource rights to different principals.) Finally, the principal who loaded a switchlet to the node and the programmer may be different entities; the switchlet may or may not include a signature from the programmer.

There are then different ways of naming a dynamic resource, each with different semantics:

- The name could be the one-way hash of the code. Assuming certain properties of the hash function, this uniquely identifies the module. The two potential drawbacks to this approach are that different versions of related services have unrelated names and that users have to discover the hash value (either through

⁸The nodes can then keep track of nonces or agent signatures that have been processed, for as long as the authorizations are valid.

⁹But remember that principals can be groups.

access to the code or by finding a trusted source that will give the user the hash value). To use the module represented by this name, a switchlet would have to trust only the module itself.

- The name could be the public key (or its one-way hash) of the module programmer, along with some other identifier assigned by the programmer (such as an ASCII string). The assumption here is that the code may be signed by the programmer (who may be different from the principal who loaded it on the active element). Version control is possible (subject to the structure of the programmer-assigned identifier). The signature would have to be verified by the active node before this name becomes available. In this case, a switchlet would need to trust the programmer before using the module represented by this name.
- The name could be the public key (or its one-way hash) of the principal who loaded the code onto the active element, along with some other identifier assigned by the principal. Since the principal had to pass an authentication/authorization check before he was allowed to load the code, there is no additional overhead imposed by this naming scheme. In this case, the switchlet must trust the installer before using the module so represented.

In fact, it is possible to combine these naming schemes as they are not mutually exclusive. Different programs may access the same resource through different names, depending on the trust policies of their respective owners. Actually accessing these services depends on the node architecture and implementation; we plan to use a portmapper-like approach, but other approaches (*e.g.*, language constructs) are possible.

As an example, imagine a principal X with a public key P who loads a new service that implements IP packet forwarding on an active node. The service was written by a programmer R who signed it with his key Q . The hash of the code is also known to be H . Any user can then access this service as:

1. $\{P, \text{"IPv4/version1"}\}$ — the IPv4 module (version 1) loaded by X ,
2. $\{Q, \text{"IPv4/version1"}\}$ — the IPv4 module (version 1) written by R ,
3. $\{H\}$ — the IPv4 module known to the user, or
4. $\{Q, \text{"IPv4/version2"}|\text{"IPv4/version1"}\}$ — the IPv4 module (version 2) if available, otherwise the previous version of the same module.

VIII. SANE IMPLEMENTATION STATUS

The SANE architecture is piece-wise implemented and the integration of the components is now underway. The AEGIS secure bootstrap architecture is currently implemented using a commercial BIOS, and has been tested up to the O.S. kernel level using the FreeBSD UNIX implementation for Intel x86 architecture machines. The AEGIS recovery algorithms are under development, but will draw on an available implementation of the IPSEC protocols for FreeBSD.

The dynamic integrity checking and availability-preservation features of the SwitchWare kernel have been implemented and tested in the prototype Active Bridge. In particular, the Active Bridge demonstrated that the use of functional languages (which are advantageous from a verification perspective) need not impose a severe performance penalty; while full details can be found in Alexander, *et al.* [34], an unoptimized prototype Active Bridge demonstrated Ethernet frame forwarding perfor-

mance of ca. 1800 frames/second and a bottleneck throughput (tested with `tcp` between two Pentia running Linux) of about 16 Mbps on 100 Mbps Ethernet connections.

Our current project is integrating these components in a SANE prototype. The methodology is to use the Utah OS Kit to build a monolithic kernel (one which exists to boot the Caml loader), and then check this using AEGIS. This project is well underway, and will provide a direct integrity chain between the low-level integrity assumptions and the running dynamic integrity checks.

IX. SANE FUTURE WORK

While the Secure Active Network Environment we have proposed provides an integrity-checked network element, secured collaboration between these elements, and a scaffolding to build a language environment capable of per-packet or per-flow restrictions and integrity checks, a great deal of work must be done to complete SANE.

We feel that two areas deserve concerted attention by members of the Active Networks research effort in general, and they are particular foci of our continuing efforts in SANE.

The first area to address is the issue of resource management. While initial active network prototyping will focus on best-effort services as a way to obtain operational infrastructure, resource management is essential to many network services such as transport of continuous media traffic. Providing explicit access to the computational and storage capabilities of a node means that there are some very difficult resource co-scheduling problems. (Examples include pinning memory, providing access to a port, scheduling a process.) An active network element must become a multiple resource multiplexer. This opens a variety of new attacks on the network infrastructure including denials of service and new covert channels. We believe that the successful approach will take the form of modern operating systems which control multiplexing at a single system layer, such as the University of Cambridge Nemesis operating system [35] or the University of Arizona Scout/Escort system [36]. These systems allow explicit resource allocation, as well as mechanism for policy enforcement. Putting services in multiplexing-controlled “containers” prevents most overload-based denial-of-service attacks.

The second issue is one of distributed programming. Our threat model has focused on building secure nodes, and providing the infrastructure upon which secure network services can be built. It is a great challenge to build systems which can examine programs, even greatly restricted programs, and decide whether or not they are safe to load. While the halting problem springs to mind, we have a much less difficult problem at the node. Even if we use a language such as the Programming Language for Active Networks (PLAN) [37] (which is restricted to provide guarantees including the termination of all programs), some programs must resort to “services” which allow a programmer (with proper authorization) to perform actions outside of the scope of PLAN itself. It is easy to imagine a well-meaning programmer writing a simple service to read a packet from an input port and write it to two output ports; with such a program a multicast facility might be constructed. If this service was indiscriminately deployed however, packets could be replicated without bound and the network could collapse of overload. This points

out the need for systematic global checking and cooperation between nodes, for which we have provided some infrastructure in SANE. The distinction this illustrates is the difference between “node safe” programs and “network safe” programs. We believe that techniques such as the Pi-calculus [38] provide a valuable avenue for exploration.

REFERENCES

- [1] J. E. van der Merwe and I. M. Leslie, “Switchlets and dynamic virtual ATM networks,” in *Proc. of the Fifth IFIP/IEEE International Symposium on Integrated Network Management*, San Diego, CA., May 1997.
- [2] D. L. Tennenhouse, J. M. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Minden, “A survey of active network research,” *IEEE Communications Magazine*, pp. 80–86, January 1997.
- [3] J. M. Smith, D. J. Farber, C. A. Gunter, S. M. Nettles, D. C. Feldmeier, and W. D. Sincoskie, “SwitchWare: Accelerating network evolution,” Tech. Rep. MS-CIS-96-38, CIS Dept. University of Pennsylvania, 1996.
- [4] Jon Postel, “INTERNET protocol,” Internet RFC 791, 1981.
- [5] Jon Postel, “Transmission control protocol,” Internet RFC 793, 1981.
- [6] D. C. Feldmeier, A. McAuley, and J. M. Smith, “Protocol boosters,” *IEEE JSAC Special Issue on Protocol Architectures for the 21st Century*, 1998.
- [7] Vern Paxson, *Measurement and Analysis of End-to-End Internet Dynamics*, Ph.D. thesis, University of California, Berkeley, 1997.
- [8] S. E. Deering, “Host extensions for IP multicasting,” Internet RFC 1112, 1989.
- [9] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “Resource ReSerVation protocol (RSVP) – version 1 functional specification,” Internet RFC 2208, 1997.
- [10] R. Atkinson, “Security architecture for the internet protocol,” RFC 1825, August 1995.
- [11] L.T. Heberlein and M. Bishop, “Attack Class: Address Spoofing,” in *Proceedings of the 19th National Information Systems Security Conference*, October 1996, pp. 371–377.
- [12] “Cert advisory ca-96.26: Denial-of-service attack via ping,” ftp://info.cert.org/pub/cert_advisories/CA-96.26.ping, October 1996.
- [13] M.D. Schroeder, “Engineering a security kernel for MULTICS,” in *Fifth Symposium on Operating Systems Principles*, November 1975, pp. 125–132.
- [14] F. Mayer M. Branstad, H. Tajalli and D. Dalva, “Access mediation in a message-passing kernel,” in *IEEE Conference on Security and Privacy*, 1989, pp. 66–71.
- [15] Dawson R. Engler, M. Frans Kaashoek, and James W. O’Toole, “The operating system kernel as a secure programmable machine,” in *Proceedings of the Sixth SIGOPS European Workshop*, September 1994, pp. 62–67.
- [16] Carl M. Ellison, Bill Frantz, Ron Rivest, and Brian M. Thomas, “Simple Public Key Certificate,” Work in Progress, April 1997.
- [17] M. Blaze, J. Feigenbaum, and J. Lacy, “Decentralized trust management,” in *Proc. of the 17th Symposium on Security and Privacy*, 1996, pp. 164–173, IEEE Computer Society Press.
- [18] Phoenix Technologies Ltd., *System BIOS for IBM PCs, Compatibles, and EISA Computers*, Addison Wesley, 2nd edition, 1991.
- [19] R. Grimes, “AT386 Protected Mode Bootstrap Loader,” </sys/i386/boot/biosboot/README.MACH>, October 1993, 2.1.5 FreeBSD.
- [20] Julian Elischer, “386 boot,” </sys/i386/boot/biosboot/README.386>, July 1996, 2.1.5 FreeBSD.
- [21] Werner Almesberger, *LILO Technical Overview*, version 19 edition, May 1996.
- [22] W. Diffie, P.C. van Oorschot, and M.J. Wiener, “Authentication and Authenticated Key Exchanges,” *Designs, Codes and Cryptography*, vol. 2, pp. 107–125, 1992.
- [23] W. Diffie and M.E. Hellman, “New Directions in Cryptography,” *IEEE Transactions on Information Theory*, vol. IT-22, no. 6, pp. 644–654, Nov 1976.
- [24] National Institute of Standards, “Digital Signature Standard,” Tech. Rep. FIPS-186, U.S. Department of Commerce, May 1994.
- [25] RSA Laboratories, *PKCS #1: RSA Encryption Standard*, version 1.5 edition, 1993, November.
- [26] Ashar Aziz, Tom Markson, and Hemma Prafullchandra, “Assigned Numbers for SKIP Protocols,” <http://skip.incog.com/spec/numbers.html>.
- [27] Li Gong and Paul Syverson, “Fail-Stop Protocols: An Approach to Designing Secure Protocols,” in *Proceedings of IFIP DCCA-5*, September 1995.
- [28] Frederick Colville Knabe, *Language Support for Mobile Agents*, Ph.D. thesis, CMU, December 1995.
- [29] Günter Kajoth, Danny B. Lange, and Mitsuru Oshima, “A security model for aglets,” *IEEE Internet Computing*, vol. 1, no. 4, July - August 1997.
- [30] Xavier Leroy, *The Caml Special Light System (Release 1.10)*, INRIA, France, November 1995.
- [31] Jon Postel, “User datagram protocol,” Internet RFC 768, 1980.
- [32] IEEE, “Media access control (mac) bridges,” Tech. Rep. ISO/IEC 10038, ISO/IEC, 1993.
- [33] R. Atkinson, “IP authentication header,” RFC 1826, August 1995.
- [34] D. S. Alexander, M. Shaw, S. M. Nettles, and J. M. Smith, “Active bridging,” in *Proc. 1997 ACM SIGCOMM Conference*, 1997.
- [35] R. Black, P. Barham, A. Donnelly, and N. Stratford, “Protocol implementation in a vertically structured operating system,” in *Proc. 22nd Annual Conference on Local Computer Networks*, 1997.
- [36] A. B. Montz, D. Mosberger, S. W. O’Malley, L. L. Peterson, T. A. Proebsting, and J. H. Hartman, “Scout: A communications-oriented operating system,” Tech. Rep., Department of Computer Science, University of Arizona, June 1994.
- [37] “Plan web page,” <http://www.cis.upenn.edu/~switchware/PLAN/>.
- [38] Robin Milner, Joachim Parrow, and David Walker, “A calculus of mobile processes, Parts I and II,” *Journal of Information and Computation*, vol. 100, pp. 1–77, Sept. 1992.

APPENDIX

I. DETECTION OF MALICIOUS NODES

In this appendix, we turn to the question of how to ensure that corrupt packets are not delivered to intermediate nodes. We take an approach which allows the originator to make a trade-off between additional computation and packet space allocated for security headers on one hand and security level on the other. Moreover, we assume that the portion of the packet which might be corrupted should not be changed by any intermediate node. Thus, if any node along the path detects any modification to the immutable part of the packet, it can determine that the packet has been corrupted (by network error or by a malicious intermediate node) and can initiate the appropriate recovery procedure. Modifications by outsiders will continue to be detected by the common authenticator shared by all nodes as described in section VI-A.

Modifications from nodes along the path (who know the secret key) are not so easily detectable. At one extreme (when using only the common authenticator), those modifications are simply undetectable. At the other extreme, including in the packet an authenticator for each node may be too resource wasteful.

A first approach would be to include a number K of authenticators, where $K < N$ and N is the number of nodes along the path. After checking the common authenticator, a node would check the list of additional authenticators for one addressed to it. If it finds such an authenticator, it verifies that one as well using the key known only to itself and the initiator. This extra verification step is relatively inexpensive, since it just verifies the common authenticator. The principal would include authenticators to a randomly chosen set of nodes for every packet.

This approach has two weaknesses. First, a malicious node can simply remove all additional authenticators. Second, such a malicious insider knows which nodes will do the checks, and therefore can modify a packet when the next node in the path is not among those, thus allowing delivery of the corrupted packet to at least one node.

The first problem can be solved by having the principal notify every node, after the path establishment, how many authenticators each packet will include. The second problem is solved by

making the authenticators anonymous: after the path is established, the initiator tells each node what its “pseudonym” will be. This pseudonym¹⁰ will be used to associate authenticators to nodes; this way, no node will know which other nodes will do the additional verification on a packet. The initiator also announces to all the nodes the list of valid pseudonyms (but not their bindings) so that a malicious insider cannot substitute an invalid pseudonym for a valid one without detection.

A malicious insider now does not know whether the next node will do the additional verification or not, since it does not know the binding between pseudonyms and nodes. It cannot remove any of the additional authenticators, since every node expects a fixed number of them on each packet. It cannot replace pseudonyms with invalid ones (since every node knows which are the valid pseudonyms) or other valid ones (since the verification would fail).

If the packet includes K additional authenticators and the path has N nodes, the probability of the next node being one that will do the additional verification is K/N in general, or $(K-1)/N$ if one of those authenticators was addressed to this node, assuming a uniform distribution of authenticators among the nodes in the path.

The initiator can decide on the value of K by balancing the level of security desired against the acceptable overheads for computation and packet size.

¹⁰This pseudonym will be a value that will be stored in the SPI field. (See Figure 6.)