

Access-Controlled Resource Discovery for Pervasive Networks

Sanjay Raman, Dwaine Clarke, Matt Burnside, Srinivas Devadas, Ronald Rivest
MIT Laboratory for Computer Science
{sraman, declarke, event, devadas, rivest}@mit.edu

Abstract

Networks of the future will be characterized by a variety of computational devices that display a level of dynamism not seen in traditional wired networks. Because of the dynamic nature of these networks, resource discovery is one of the fundamental problems that must be faced. While resource discovery systems are not a novel concept, securing these systems in an efficient and scalable way is challenging. This paper describes the design and implementation of an architecture for access-controlled resource discovery. This system achieves this goal by integrating access control with the Intentional Naming System (INS), a resource discovery and service location system. The integration is scalable, efficient, and fits well within a proxy-based security framework designed for dynamic networks. We provide performance experiments that show how our solution outperforms existing schemes. The result is a system that provides secure, access-controlled resource discovery that can scale to large numbers of resources and users.

1 Introduction

Resource discovery is one of the fundamental challenges that must be faced in the context of pervasive computing. Simply stated, the goal of resource discovery is to provide a user with a snapshot of the computational environment in which he is operating. Resource discovery is vital to enabling operation in pervasive networks as the network state is unpredictable. The dynamism of pervasive networks also brings rise to problems of security. Although often overlooked, security is a critical component necessary for the practical realization of pervasive computing. As a resource provider, we want to guarantee that foreign users that enter our environment will not be able to act maliciously. Similarly, as a user in a foreign environment, we want to know what resources we are able to use and which ones we can trust. Such access restrictions are easily handled in fixed networks as foreign users can simply be denied admission to the

network. But the fundamental notion behind pervasive computing gives rise to the idea of resources and users of varying privileges interacting in the same environment. Further complicating the issue is that pervasive computing environments handle a diverse and heterogeneous set of users and resources [2], including computationally-limited devices, so it is important to enforce a security framework that can be extended to many disparate resources. Communication channels must be secure and access control must be granted to resources in order to regulate usage. While several systems [7, 15, 9] propose resource discovery solutions for dynamic environments, they do not consider how the integration of security protocols influences scalability and performance.

Resource discovery systems are typically implemented in the network layer, below security, allowing networks to overlay any desired security protocol. An access control framework can be layered over a resource discovery protocol, but these two protocols seem to have different goals. The goal of a resource discovery system is to find the resource or service that best matches the criteria for which a user is looking. On the other hand, a security protocol that enforces access control is concerned primarily with allowing users to perform authorized operations on protected resources. The problem is that the best criteria-matching resource (e.g. “the nearest, least-loaded printer”) may not necessarily be a resource to which a user has access.

The primary focus of this paper is to address the issue of resource discovery in a pervasive computing environment. More specifically, this paper presents a system that integrates access control with resource discovery in order to enable scalable and efficient operation. This paper describes a resource discovery system that is scalable and efficient and is designed to elegantly integrate with a proxy-based security system [3].

The proxy-based security system uses a distributed SPKI/SDSI protocol [13] which allows for private, encrypted communication between heterogeneous lightweight devices in a pervasive computing

environment.

The architecture presented in this paper makes four key contributions:

- A scalable model for resource discovery based on the Intentional Naming System [1] that integrates access-control information with service information.
- Integration of access-controlled resource discovery with a proxy-based security infrastructure to provide secure and authentic communication in a pervasive computing environment.
- Implementation of a resource lookup algorithm that makes access control decisions while finding the best resource.
- Design of lightweight, high-performing access control lists.

In the remainder of this section, we briefly outline the proxy-based security architecture that is a basis for our system. We summarize the resource discovery problem in terms of a simple scenario in Section 2. Section 3 details our system architecture and describes how we have developed an access-controlled resource discovery system. We present the advantages and performance evaluation of our system in Section 4. We discuss some related systems in Section 5 and conclude the paper in Section 6.

1.1 Background

The resource discovery system presented here is designed to be an integral part of a larger proxy-based security architecture [3]. Resources are defined to be any piece of hardware or software that is providing a service to members of the network. A resource may be location-aware. In this architecture, each resource has an associated trusted software proxy. A proxy is software that runs on a network-visible computer and its primary function is to execute commands on behalf of the resource it represents. Proxies store certificates and other information for the resource they represent and are trusted implicitly. Proxies communicate with each other using a protocol based on SPKI/SDSI (proxy-proxy protocol). A separate resource-proxy protocol is used for secure communication between resources and proxies. Having two different protocols allows us to run a computationally-inexpensive security protocol on impoverished resources and a sophisticated protocol for resource authorization on more powerful resources. Figure 1 shows an overview of this system and the protocols it uses.

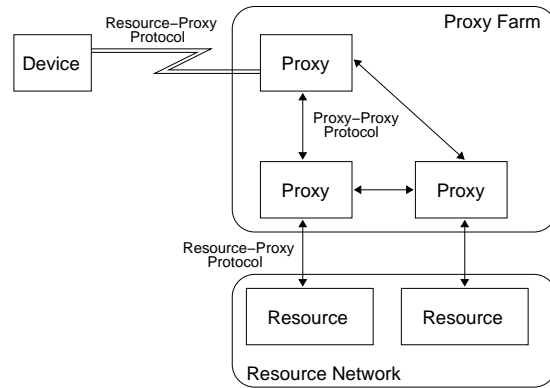


Figure 1: An overview of the basic components in the proxy-based security infrastructure. Proxies communicate via the proxy-proxy protocol. Devices and resources communicate with proxies via the resource-proxy protocol.

The proxy-proxy protocol layers SPKI/SDSI access control over an application protocol, which in turn is layered over a key-exchange protocol. This allows us to deal with a variety of application protocols that may be implemented across wired or wireless links in a heterogeneous network. SPKI/SDSI features an elegant model for access control lists (ACLs) and delegation of authority.

2 The Problem Restated

The problem that the system in this paper solves is that of how to scale a system of resources that are protected by access control. It is tremendously inefficient if a user repeatedly attempts to contact a resource that he is prohibited from using. One only has to consider an environment with a large number of protected resources. If a user has no knowledge of which resources he can access, it could take an exhaustive computational effort to find an accessible resource. In order to gain scalability and efficiency, the resource discovery system needs to know about access control privileges so that it can return the best resource to which *a user has access*. By knowing the user's authorizations (i.e., the groups to which he has membership) and the access control lists of the supported resources, a resource discovery system can effectively meet this goal.

2.1 A Simple Scenario

This issue is especially pertinent when dealing with networks where the state of the network is highly dy-

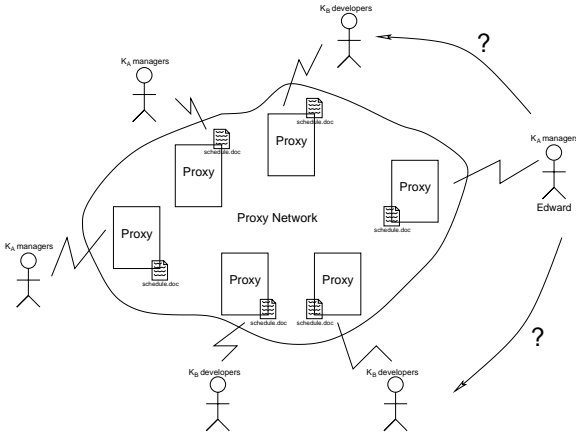


Figure 2: This figure illustrates the conflict experienced by a resource discovery system in an access-controlled environment. How does Edward find the closest, accessible copy of `schedule.doc` without performing an exhaustive search?

namic. It is very plausible that a user will not know exactly what resources are available, nor will the user know which he is authorized to use. As a simple example, consider an environment which treats all devices in the network as resources in a peer-to-peer application. Figure 2 illustrates the following scenario:

Edward, a manager at a large software firm, arrives in the morning at a conference with his location-aware device. Upon arriving and coming online, Edward wants to download his personalized conference schedule for the given day. At this conference, there are two tracks: one for managers and one for software developers. Thus, the users in the system are divided into two groups, K_A managers and K_B developers. All the users already at the conference have a document, `schedule.doc`, in their repository, but the document is track-specific. That is, the copy of `schedule.doc` that members of K_A managers have is different than the copy that members of K_B developers have. When Edward comes online, he wants to synchronize his copy of `schedule.doc` by getting the latest version. The conference is spread out over several buildings and the users are spatially far apart. Because the physical area of the conference is large, there is no central repository for the schedules. Instead, schedule distribution and synchronization happen peer-to-peer. As a member of K_A managers, Edward must get the doc-

ument from another member of *his group*. Members of K_A managers do not have access to the schedules of members of K_B developers, and vice versa. Edward would also like to get the schedule from the geographically closest user, in order to minimize his delay and make the synchronization process as fast as possible.

2.2 Problems

This scenario creates a conflict of interests. Not only must Edward find the closest user, but also must find a user that is in his group (a resource to which he has access). A simple resource discovery system could easily tell Edward the location and identity of the closest user. This problem has been solved many different ways [1, 16, 11]. But, how does Edward know if the physically-closest user is a member of his group? And, if this user is not a member of his group, where exactly is the closest member of Edward’s group? Mobility of the users only further complicates the issue. It would be time consuming and inefficient for Edward to blindly search for the closest member of his group. The only way in which a resource discovery system can identify the closest, accessible resource is to know ahead of time Edward’s identity and authorizations.

2.3 A Naïve Solution

Before presenting our solution, it is instructive to outline a naïve solution. This solution will be used as a baseline of comparison in terms of performance and will be important for the analysis of Section 4. Resource discovery systems that do not incorporate the ideas presented in this paper will typically operate by returning the address(es) of the best criteria-matching resource. It will become clear that issues of scalability and efficiency are major obstacles with such a system.

In attempting to discover the geographically-closest user, Edward will query the resource discovery system through his personal proxy. The proxy will tell the resource discovery system to “find me the closest user”. Ideally, Edward would like to contact the closest, *accessible* user, but this resource discovery system does not know anything about Edward’s identity or authorizations. In response to the query, the resource discovery system will return a list of the geographically-closest users to Edward’s proxy. At this point, Edward’s proxy does not know which of the resources in the list are accessible to him. The only reasonable way for the proxy to proceed is to

sequentially iterate through the resources in the list in the hope that they are accessible. The proxy must engage in some sort of authorization check in order to determine if the user has access to the resource. As long as a contacted resource fails, the proxy will have to repeat the process.

This approach can be inefficient and surely is not scalable. If a given user has access to every resource in the network, then the efficiency of access control is not an issue. But, in most heterogeneous environments, users are assumed to be diverse and access privileges will exhibit the same differences. In Edward’s scenario, if he is not close to any users of his group, he would have to iterate through many inaccessible resources before finally finding a match. Edward is faced with executing a process on the order of $O(n)$ if there are n other resources in the network. The results of Section 4 will illustrate this point.

3 System Architecture

A better approach would be to give the resource discovery system knowledge about the access control lists that protect the resources. We require that the designed system be secure, efficient, scalable, and robust. In order to meet our goals, the Intentional Naming System (INS) [1] was selected. The solution presented here uses several modifications to intentional naming that enables access control decisions to be made while finding the best resource. Before detailing our solution, we summarize INS as a standalone resource discovery system.

3.1 Intentional Naming Overview

Intentional Naming System (INS) is a resource discovery and service location system intended for dynamic networks. INS is ideal for dynamic networks because an application only needs to tell the service the resource characteristics it is seeking. Since the availability of resources may be dynamic, these systems require a naming service that is just as flexible. The Domain Naming Service (DNS) works well for static networks since an application can be fairly confident in the names of resources. INS provides users with a layer of abstraction so that applications do not need to know the availability or exact name of the resource for which they are looking. A simple example of a user’s request in INS is to find the nearest, least-loaded printer. DNS would require the user to know the exact name of the resource, such as `pulp.lcs.mit.edu`.

INS uses a simple language based on expressions called *name specifiers*, which are composed of an

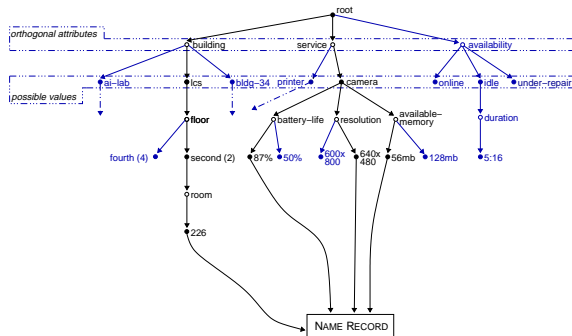


Figure 3: A graphical view of an example INS name-tree. The name-tree consists of alternating layers of attribute-nodes, which contain attributes and value-nodes (possible values). Value-nodes contain pointers to all the name records they correspond to. The bold region shows an example name-specifier.

attribute and value. An attribute is simply a category by which a resource can be classified. For example, a camera in the system can be described by its **resolution**, **battery-life**, and/or **available-memory**. An INS name, or *intentional name*, is a hierarchy of these atomic name specifiers. An example of an INS name is `[service=camera [resolution=640x480] [battery-life=87%] [available-memory=56mb]]` to describe a camera with the specified properties.

INS is comprised of a network of *Intentional Name Resolvers (INRs)* that serve client requests for resources and maintain information about the searchable metadata of each resource. Data is represented in the form of a dynamic *name-tree*, which is a data structure used to store the correspondence between name specifiers and the destination resource. The structure of a name-tree strongly resembles the hierarchy of a name specifier. Name-trees consist of alternating levels of attributes and values, with multiple values possible at each attribute. A particular name specifier is resolved by traversing the tree, making sure to visit all the corresponding attribute-value pairs of the target resource. Each leaf value in the name-tree has a pointer to a *name-record*, which holds the physical location of the resource. Figure 3 illustrates an example name-tree.

3.2 Security Integration with INS

The solution presented here uses several modifications to intentional naming that enable access control decisions to be made while finding the best resource. While INS does allow for a security framework to be

layered over it, we have already seen how a system can benefit from integrating access control decisions with resource discovery. INS is extended in the following three ways to provide access-controlled resource discovery:

1. implementation of a real-time maintenance of the access control lists in the INS name resolvers,
2. introduction of a certificate-based authorization step during resolution of an INS request, and
3. design of a lookup algorithm that prunes the possible name records by eliminating resources based on a user's identity and authorizations.

In the following sections, the key extensions of INS are presented. Finally, we will return to the scenario that is discussed above to see how this new system integrates access-control information and INS knowledge to efficiently return the best, accessible resource. Another key factor influencing this design was its inclusion as a small piece of a larger security infrastructure. Therefore, some of the components of this design were chosen to leverage the existing functionality of this proxy-based security system.

3.2.1 Storage of ACLs in INS

Assuming that resources have the ability to inform INS of the access control lists that protect them, how can these lists be properly stored in the INS knowledge base so that they can be referenced when making resource decisions? INS uses a name-tree to store its knowledge about resources in the system. Name-trees are dynamic, changing their structure based on how resources advertise and re-advertise themselves to INS.

An access control list is treated as an additional attribute that defines a resource. If one can specify a **camera** based on its **resolution**, an access control list is just another way to classify the camera. In order to store ACLs as attribute-value pairs, a new type of attribute was introduced. Previously, all attributes were treated as *searchable*, in that they were used as a dimension along which a resource can be explicitly queried. But, when a user makes a request for a resource, the user cannot specify the ACL attribute-value pair in the query. Nor do we want the ACLs being represented as additional branches in the name-tree. So, in order to store ACLs, the concept of a *hidden* attribute was defined. INS attributes are now defined as searchable or hidden, with the only hidden attribute being that of the ACL. When advertising its service profile, a resource will advertise its ACL like any other searchable attribute, but the

name resolvers are responsible for denoting the ACL as a hidden attribute and storing it on the name-record for the particular resource.

Storing ACLs as attribute-value pairs is advantageous because we do not change the manner in which data is stored and we do not have to radically alter the way in which queries are handled (ref. Section 3.2.2). The structure of the name-tree remains the same, while the hidden attributes are stored directly on the name-records for each resource.

3.2.2 Redesign of lookup algorithm and ACL propagation

[1] describes the LOOKUP-NAME algorithm that INS uses to retrieve name-records for a given name-specifier. This algorithm operates by pruning attribute branches of the name-tree that fail to match the given search criteria, ultimately arriving at a subset of all the name-records that contains the possible matching resources. This algorithm works well with the way name-trees are organized in INS. Since the name-trees consist of alternating levels of attributes and values, it is very easy to prune branches of the tree while progressing through the target name-specifier. But, left alone, this algorithm fails to work with hidden attributes such as ACLs.

Due to the transparency that is required, users will not explicitly construct queries with ACL name-specifiers. One option for determining a user's accessible resources would be as follows. First, run the LOOKUP-NAME algorithm to completion, arriving at a list of criteria-matching resources. At this point, INS would have a handle to the name-records for each of the matching resources. We could proceed by iterating through these possible name-records and checking whether the user making the request is on the ACL. While this approach will save us considerably over the approach of contacting each of the resources for access decisions, it still is inefficient. A closer inspection of the LOOKUP-NAME algorithm reveals additional ways in which this process can be optimized.

We designed a modified algorithm, LOOKUP-NAME-AC, that eliminates potential name-records *while* pruning S , the set of all possible name-records. The LOOKUP-NAME-AC algorithm operates under some assumptions on the state of the INS name-tree. In order for the algorithm to terminate successfully, the algorithm assumes that each value node in the name-tree contains an intermediate ACL. This intermediate ACL is computed to be the logical OR (\vee) of the intermediate ACLs stored at all of the value nodes that are its children in the INS name-tree. Be-

ginning at the value-nodes that contain pointers to name-records, intermediate ACLs are computed. For these leaf nodes, the intermediate ACL is simply the ACL of the name-record to which it points. After computing the ACLs at these leaf nodes, the intermediate ACLs for the parent nodes are computed all the way up the name-tree. OR'ing (\vee) multiple access control lists happens at the "entry" level. That is, the result of the logical OR (\vee) of two ACLs is a new ACL with every entry that exists in either of the two ACLs. For example, if $acl_a = [e_1, e_2, e_3]$ and $acl_b = [e_1, e_2, e_4, e_5]$, then:

$$acl_a \vee acl_b = [e_1, e_2, e_3, e_4, e_5], \quad (1)$$

where the notation $acl = [e_1, \dots, e_n]$ indicates that e_1, \dots, e_n are entries of the ACL. Figure 4 illustrates how ACLs are propagated up the INS name-tree from the leaf nodes.

The modified algorithm is similar to its predecessor, except now it eliminates candidate records based on whether the user is included in intermediate ACLs. This new algorithm takes the user's identity and authorization rules as arguments. For each name-specifier in the INS query, INS will prune branches that do not match the search criteria and that do not contain the user in their intermediate ACLs through a series of recursive calls. When the algorithm terminates, S will only contain the relevant, accessible name-records. By taking the OR of the ACLs, we enable access control decisions to be made while INS is locating the proper name-record, eliminating the need to iterate through inaccessible resources and branches of the tree. This simplifies the task of the lookup algorithm as well as potentially reducing the amount of the name-tree that needs to be traversed. This algorithm terminates without the need to backtrack and does not ever check a given ACL more than once. The additional cost of this algorithm, though, is clearly in these checks that the algorithm must make for each name-specifier, but we argue in Section 4 that this tradeoff is still advantageous.

3.2.3 Dynamic maintenance of name-trees

ACLs are resource properties that may change. Groups or keys may be added or removed, or the operations allowed by a particular group/key may be changed. In dealing with name-tree maintenance, there are three qualities that any design should achieve:

- **Freshness.** Our primary goal is to keep the state information in the INS name-tree *fresh*. At

any point in time, we want to know with high probability that the access control and resource information is up-to-date.

- **Responsiveness.** The maintenance procedure should be *responsive* to changes made to the access control information. It is important that changes to ACLs are rapidly reflected in the INS knowledge base.
- **Authentication and Privacy.** Finally, maintenance updates should be *authentic* and *private*. There are a whole suite of attacks that can be centered around unauthentic updates (replay, DoS). Also, entities should not be able to maliciously learn sensitive information about resources. For these reasons, the security of maintenance updates is very important.

Many of these issues have been considered when designing INS for service updates, so our focus is specifically on how access control updates are handled. Responsiveness is achieved by using triggered updates which are fired when an ACL changes state. Periodic updates are also used to enforce freshness. The utility of these updates comes from the fact that ACLs typically have expiration times. Clearly, the update period should be chosen such that it is less than the ACL expiration time ($T_{update} < t_{expire}$) but not so small that it unnecessarily floods the network with update packets. Upon receiving an update request, INS actively modifies its name-tree to reflect the current state of access rights and intermediate ACLs are recomputed. Handling the privacy and authenticity of these messages, as well as the authenticity of messages in which a resource updates INS with its other attributes, is a subject of ongoing research.

3.2.4 User authorization rules

In order for this system to function, INS needs access to the user's set of current authorizations. The modified lookup algorithm depends on knowing the user's identity and the groups of which he is a member. Each proxy in the system stores a user's signed SPKI/SDSI certificates. [5] describes an efficient algorithm for determining, from a set of SPKI/SDSI certificates, the access control groups that a particular user is a member of and the operations that he is allowed to perform. A complete and detailed description of these procedures is found in [5], but this is well beyond the scope of this paper. In essence, a (finite) transitive closure is taken over the certificates, and rules representing the user's authorizations are extracted. The rules are simple and not signed.

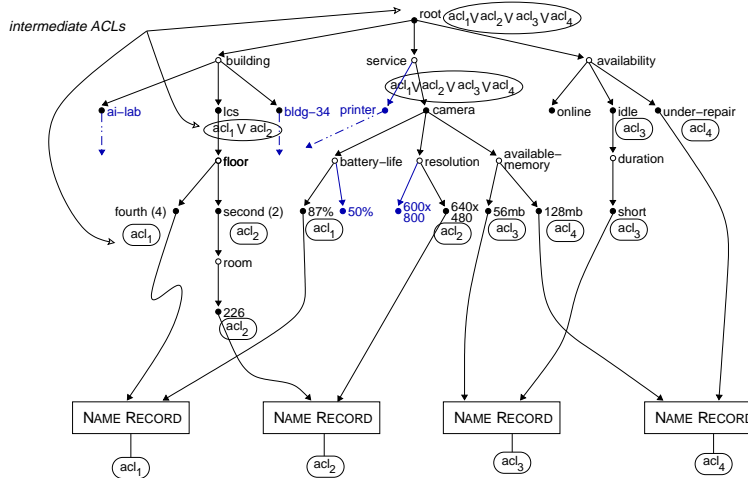


Figure 4: This shows how ACLs are propagated from the leaf nodes up the INS tree to the root of the data structure. At each intermediate value-node in the tree, an ACL is stored and is computed by taking the logical OR (\vee) of the ACLs at all of the child nodes.

However, each rule has a representation as a signed user certificate, or a chain of signed user certificates. The closure algorithm is run when there is a change in the user's certificates, such as when he acquires a new certificate, or when one of his certificates expires.

The proxy presents the user's authorization rules to INS with the user's query. INS uses the rules to check if the user is on an (intermediate or leaf) ACL contained at a node in the INS tree (using the LOOKUP-NAME-AC algorithm). An important point is that these ACL checks performed by INS can be made fast and efficient. The ACL check is used to determine *if* a user is on an ACL, and it is not necessary for INS to know the proof that the user can generate to show that he is on the ACL.

When INS has completed its searching and returned an address, the proxy will then use a secure authentication and authorization protocol to contact the resource. The modified INS system we present now returns only resources to which the user has access, so the proxy should only have to execute this security protocol once.

3.3 The Scenario Revisited

After presenting the components of the access-controlled resource discovery system, it is helpful to revisit the scenario presented in Section 2.1 to see how this new system handles the same problem.

Again, Edward is looking to obtain a copy of his schedule, `schedule.doc`, from the closest user in his group. Edward places a request for the document via his proxy. Edward does not explicitly have to

indicate to his proxy his group membership or the fact that he wants to retrieve the document from another group member; this is handled automatically by his proxy. Edward's proxy contacts an INR with which it has previously registered. It then queries the INR for the best accessible resource, translating the request specified by Edward to an INS-specific name-specifier. Edward's proxy also computes his authorization rules (they may be computed on the fly or pulled from the proxy's cache) and sends them along with the request to the INR. The INR, which has received access control advertisements from all the registered resources in the network, takes the request and the user's authorizations and executes the LOOKUP-NAME-AC algorithm. After a single execution of this algorithm, the INR returns the closest, accessible resource to Edward's proxy. Edward's proxy then uses a secure protocol to contact the resource and uses a standard secure copy protocol to retrieve the file from the resource. Because INS knew about Edward's group membership, it returned a resource that is accessible, meaning the time-consuming security protocol *would only have to be executed once*.

4 Evaluation

A prototype system was implemented in Java using INS 2.0, a pure Java implementation of INS. In this section, a formal evaluation of this system is presented. The overall goal is to quantify how our design outperforms a resource discovery system that does not integrate access control with resource decisions.

These experiments were all conducted using off-the-shelf Intel Pentium II 266MHz computers with a 512 KB cache and 128 MB RAM, running Windows NT Server 4.0. The software was built and run using Sun’s Java Virtual Machine version 1.3.

4.1 Comparison of resource retrieval time

A measure of the time savings of our solution is necessary to evaluate its effectiveness. As a baseline, this system will be compared to a basic scheme, where INS is used as the resource discovery system, but *does not* have access to ACLs or the authorizations of the requester. This basic scheme was described in detail in Section 2.3. For convention, we will assume that the user, U is operating in a network with n total resources.

To understand the performance gains of this new solution, we must analyze the time it takes U to successfully access the most optimal resource and compare this time in both the *basic* and *access-controlled* systems. This time is denoted as t_{BASIC} for the basic scheme and as t_{AC} for the access-controlled system. Each of these time values can be generally expressed by the following equation:

$$t_x = t_{query} + \left(\sum_{k=1}^n b_k \cdot (t_{latency} + t_{acl-check}) \right) + t_{crypto} \quad (2)$$

t_{query} is the *query time*, the time it takes the resource discovery system to respond to U ’s request. t_{query} also includes any time U ’s proxy uses to prepare the request. b_k is a boolean value that is 1 if U contacts resource k and 0 if U does not. $t_{latency}$ is the round-trip network latency between two proxies. This is essentially the time it takes U to retrieve a resource’s ACL over the network. $t_{acl-check}$ is the *ACL-check time*, the time it takes for a simple ACL check to be performed. ACL checks were made very fast with our adopted implementation (as will be shown later in this section). Finally, t_{crypto} is the time it takes U to derive the full authorization proof and for this proof to be verified by a particular resource’s proxy.

4.1.1 t_{BASIC}

In the basic scheme, the time for U to successfully access the most optimal resource is given by the following equation:

$$t_{BASIC} = t_{query_{BASIC}} + \frac{1}{p} \cdot (t_{latency} + t_{acl-check}) + t_{crypto} \quad (3)$$

This derivation of t_{BASIC} can be found in [12]. $t_{query_{BASIC}}$ is the time it takes the LOOKUP-NAME algorithm to execute and p is the probability U has access to a given resource.

4.1.2 t_{AC}

Similarly, the time to retrieve a resource using our access-controlled solution is given by:

$$t_{AC} = t_{query_{BASIC}} + D_n \cdot t_{acl-check} + t_{latency} + t_{crypto} \quad (4)$$

The key difference is that t_{AC} is not dependent on the likelihood that U has access to a given resource. Instead, the query time, $t_{query_{AC}}$, depends on D_n , which represents the number of ACL checks that will have to be made while traversing the INS name-tree. It is a function of the number of resources in the network (n), but also is affected by the complexity of the name-tree and name-specifiers. For more details, see [12].

4.1.3 Name Lookup Performance, $t_{query_{BASIC}}$ and $t_{query_{AC}}$

To quantify the difference between $t_{query_{BASIC}}$ and $t_{query_{AC}}$, we constructed a large, random name-tree and timed how long it took the tree to perform 1000 random lookups using each algorithm. The name-tree and name-specifiers were chosen uniformly according to the parameters defined in [1] ($r_a = 3$, $r_v = 3$, $n_a = 2$, and $d = 3$). n , the number of distinct, unique names in the tree, was varied from 1 to 13000 in increments of 100 to see how $t_{query_{BASIC}}$ and $t_{query_{AC}}$ vary with increasingly large name-trees. The maximum heap size of the JVM was limited to 64MB, thus limiting the range of the experimentation.

Figure 5 shows the results of this experiment. Using the basic LOOKUP-NAME algorithm, the performance went from a maximum of around 700 name lookups/sec to a minimum of 200 lookups/sec. From Figure 5, it is evident that as the number of names in the name-tree increases, the lookup rate decreases. As a result, the amount of time required for a single lookup increases. But, the drop-off is not as drastic as one would think and clearly is not linear. For a moderately large system with approximately 2000 resources (or names), the average lookup time is around




Figure 5: The lookup rate (lookups/sec) is plotted against the number of names in the name-tree. As the number of names increases, the lookup rate progressively gets smaller, starting from a maximum of around 700 lookups/sec to 200 lookups/sec.

1.8 ms. For small systems on the order of hundreds of resources, the lookup time is around 1.4 ms. These times are small and the difference in lookup times between the small and large systems is minimal.

The experiment was repeated in the access-controlled case. Each resource was initialized with ACLs containing 10 unique entries and the intermediate ACLs were computed. Figure 6 presents the performance results of the LOOKUP-NAME-AC algorithm as the number of names in the tree varied from 1 to 3500. As is evident from this figure, the lookup rate is significantly reduced from the rate without the ACL checks. The experiment was terminated at a maximum of 3500 names due to memory constraints of the JVM. With approximately 100 name-records in the tree, a rate of 325 lookups/second was achieved. In the non-access-controlled case, this rate was much higher at around 700 lookups/sec. At approximately 3500 name-records, the rate of the LOOKUP-NAME-AC algorithm was at 240 lookups/sec, indicating only a drop of in about 90 lookups/sec. Conversely, the rate in the basic case dropped to 450 lookups/sec with 3500 names, indicating a drop of 250 lookups/sec.

Table 1 details the average lookup times for the two algorithms for varying sizes of the name-tree. The difference between the lookup times is on the order of few milliseconds and can be attributed directly to the intermediate ACL checks that are made. In the following section, it will be shown that $t_{acl-check}$, the time for a simple ACL check is on the order of approximately .07 ms. Based on the name-trees we used during the experimentation, we can calculate approximately 15 intermediate ACL checks. This roughly




Figure 6: The lookup rate (lookups/sec) is plotted against the number of names in the name-tree. Each name in the name-tree is protected by an ACL with 10 unique keys. As the number of names increases, the lookup rate progressively gets smaller, starting from a maximum of around 325 lookups/sec to 240 lookups/sec.

Names in Name-Tree	Average Lookup Time (ms)	
	LOOKUP-NAME-AC	LOOKUP-NAME
100	3.24	1.45
200	3.23	1.47
500	3.35	1.48
1000	3.52	1.61
1500	3.66	1.76
2000	3.80	1.88
2500	3.94	2.04
3500	4.23	2.31

Table 1: This table shows the average lookup time experienced by the two algorithms for varying sizes of the name-tree.

accounts for about a $15 \times .07 \approx 1.05$ ms difference between the lookup times. The numbers in Table 1 seem to support this back-of-the-envelope calculation.

4.1.4 Access Control List Performance, $t_{acl-check}$

One of the fundamental differences between a basic solution and ours is the use of ACL checks during the name-lookup process. In order to determine the cost of an ACL check, random large ACLs were constructed with the number of distinct entries in the ACL ranging from 1 to 14000 and the number of ACL checks that could be executed in the span of a second was measured. Figure 7 illustrates the results of this experiment. As expected, as the number of entries in the ACL grows, the ACL check rate decreases




Figure 7: The ACL check rate (in ACL checks/sec) is plotted against the number of entries in the ACL. It is evident that the rate decreases with an increasing number of entries, but only slightly. Also of note, are the four regions of concentration of points.

logarithmically. ACLs in our system are represented by red-black trees (binary trees), keyed by the users’ public keys, that guarantee a $\log(n)$ time cost for adding new indices and looking up values. As the number of entries in the ACL goes from 1 to 1000, the check rate decreases by 500 checks/sec. A similar rate decrease can be seen as the number of entries is varied from 1000 to 10000.

Figure 7 shows five stratified regions of lookup rates, that correspond to the number of decisions that must be made in order to find a key in the ACL. Depending on where a key is located in the range of possible keys, the number of decisions to find it in the tree can vary. For an ACL of 1000 entries, the time it takes to perform an ACL check can be one of the following values: .083 ms, .074 ms, .067ms, or .061ms (according to the four different regions in the graph). These values are an order of magnitude smaller than the time taken by the LOOKUP-NAME algorithm to find a name. Therefore, the idea of making several ACL checks during the name retrieval process adds a minimal time cost and seems very reasonable.

4.1.5 Round-Trip Network Latency, $t_{latency}$

$t_{latency}$ is the round-trip network latency between proxies in the network. It is a fundamental component of the resource retrieval time in the basic solution (t_{BASIC}), which requires a client proxy to explicitly contact potential target proxies in order to determine access privileges. To estimate this param-




Figure 8: The results of the simulation are shown here. The RTT of packets sent between proxies is constant throughout the packet flow. The mean RTT between two network proxies is 48.37 ms.

eter, simulations were run in ns [8].

A precise measure of $t_{latency}$ is somewhat subjective, as the exact value of the round-trip time between two proxies depends on the network infrastructure, number of hops between proxies, current traffic conditions, link bandwidth, and any additional network characteristics. For simulation purposes, we adopt a network structure where proxy-proxy communication will take at most two hops. Two routers are linked together, with each router containing seven end proxies, each sending packets according to the following traffic flows. A third router is connected to INS and the other two routers. Therefore, in order to communicate with another proxy, a proxy must only send packets through two hops. The links between proxies and routers each have a bandwidth of 133 Mbps and a propagation delay of 5 ms. The router-router links have a bandwidth of 100 Mbps. There are three main traffic flows in this network, namely Proxy-Proxy traffic, Proxy-INS service updates, and Proxy-INS requests.

A single proxy-proxy flow was started between two proxies and the round-trip time for each packet was measured over a span of thirty seconds. Figure 8 shows the results of this experiment.

As is evident from this figure, the round-trip time stays almost constant throughout the duration of the traffic flow. Initially, there is some variance as TCP uses a slow-start mechanism to find the optimal window size. But, after equilibrium is reached, the mean RTT of proxy-proxy communication is 48.37 ms. It is worth noting that in a network with many resources, this number is a best-case scenario. The link bandwidths used were large, the propagation delays were

small, and the two-hop assumption will break down as the number of resources increases. Despite using favorable conditions, we see that $t_{latency}$ is three full orders of magnitude larger than $t_{acl-check}$. As will be shown in Section 4.1.6, this result plays a key role in determining the efficiency of our access-controlled resource discovery system.

4.1.6 t_{AC} versus t_{BASIC}

In this section, we analyze the difference in retrieval times between the two solutions. Subtracting Equation 4 from 3, we get:

$$\begin{aligned} \Delta_t(n) &= t_{BASIC}(n) - t_{AC}(n) \\ &= \frac{1}{p}(t_{latency} + t_{acl-check}) - \\ &\quad (D_n \cdot t_{acl-check} + t_{latency}) \end{aligned} \quad (5)$$

From Equation 5, we can see that whether the access-controlled scheme outperforms the basic scheme depends on whether $\frac{1}{p} \cdot (t_{latency} + t_{acl-check})$ is greater than $(D_n \cdot t_{acl-check} + t_{latency})$. If it is, we can conclude it is more efficient for INS to perform the ACL checks as it descends down its name tree, rather than leaving this up to the user’s proxy. In order to make this comparison, we consider our scenario (in Section 2.1) with 1000 total users divided equally among the two groups (K_A **managers** and K_B **developers**). Therefore, the probability that Edward has access to any given resource is $p = 0.5$. If we also assume the structure of the name-tree is as described previously, $D_n = 15$. From our experiments in Section 4.1.4, we will assume an ACL check with 1000 entries per ACL takes .083 ms. Finally, the latency between proxies will be assumed to be 48.37 ms (as calculated in Section 4.1.5). Using these parameters, the difference in lookup time is:

$$\begin{aligned} \Delta_t(n) &= \frac{1}{0.5}(48.37 + 0.083) - (15 \cdot 0.083 + 48.37) \\ &= 47.291 \text{ ms} \end{aligned} \quad (6)$$

Even with the parameters chosen to favor the basic solution, the access-controlled solution wins by a large margin. It is likely that this is a conservative estimate. With 1000 resources in the network, $t_{latency}$ will likely be greater than 48.4 ms as the propagation delays of the links will increase and the number of hops between proxies will increase. Furthermore, if p becomes smaller, the basic solution is subject to more trips across the network, making our savings greater. The main difference in the resource retrieval times for

each solution can be attributed directly to the fact that ACL checks are extremely fast. Our solution is not subject to the network latency and the three orders of magnitude saved in performing an ACL check give our solution a clear advantage. The query time saved in the basic solution is minimal compared to the time that the ACL checks save.

It is also useful to estimate the actual time it takes to retrieve a resource with each solution. With 1000 resources in the system, $t_{query_{BASIC}}$ is 1.61 ms and $t_{query_{AC}}$ is 3.52 ms (taken from Table 1). t_{crypto} can be estimated to be 330 ms for a user with 2 certificates in his cache ([4] details this derivation). Therefore,

$$\begin{aligned} t_{AC} &= 3.52 + 48.37 + 0.083 + 330 \\ &= 381.97 \text{ ms} \end{aligned} \quad (7)$$

$$\begin{aligned} t_{BASIC} &= 1.61 + \frac{1}{0.5} \cdot (48.37 + 0.083) + 330 \\ &= 428.52 \text{ ms} \end{aligned} \quad (8)$$

Overall, the savings of our solution in terms of resource retrieval time are very significant.

4.2 Tradeoffs

In this section, we have seen how our access-controlled resource discovery system compares to basic solutions that do not integrate access control. The premise our solution makes is that basic solutions scale poorly and are based on inefficiencies that limit the performance of the system. Specifically, finding a resource requires explicit contact to check access privileges. From the experiments, we have verified this by showing our solution significantly reduces the resource retrieval time. At the same time, it makes a large system with many resources manageable and efficient. A similarly-sized system may be inoperable under the basic resource discovery approach.

While saving time, our solution does add greater requirements for storage to INS. Our solution is based on storing ACLs in the name-tree, something that basic solutions do not need to do. In [12], we have shown that in order to store intermediate ACLs in the name-tree, we require, on average, 3.75 times more storage capacity than the normal INS tree. But, storage is cheap and can be solved simply by adding more memory to each INR. On the other hand, saving time is not as simple as installing additional components to each router. As such, this tradeoff is one that is worth making.

Secondly, our system requires more maintenance and introduces new synchronization issues. ACLs

that are stored in the name-trees must be correct and fresh. Anytime a service changes, the ACLs must be recomputed. In [12], we experimentally show that a name-tree with 1000 resources takes approximately 3.7 seconds to fully propagate its ACLs. This time is two full orders of magnitude greater than average lookup times, meaning that it is a relatively time consuming process. But, we argue that this is compensated by the fact that we expect ACL updates to be very infrequent. While this time is large, this becomes a common scheduling problem and can be solved easily by existing techniques.

5 Related Work

There are several protocols that have been developed that provide resource discovery services. This section presents some of these protocols and gives a brief analysis of each.

Jini network technology [14] is a Java environment developed by Sun Microsystems that supports resource discovery. The overall goal of Jini is to turn a network into a flexible, easily-administered tool with which resources can be easily found by clients. In practice, Jini extends the Java application environment from a single virtual machine to a network of machines. Communication between a network of virtual machines occurs by exchanging serialized Java objects over Java Remote Method Invocation (RMI). Jini provides a lookup service to all clients in the network. The lookup service enables clients to query for the resources through a standard Java interface. While Jini offers a great deal of flexibility, its reliance on the Java Virtual Machine (JVM) makes it only as secure as the minimally secure JVM. The Jini architecture does not include any security in addition to the normal Java security facilities (for protecting the client JVM from malicious code), and the security aspects of RMI are insufficient for a trust-based security model [6]. RMI execution is layered on top of Java sockets and is abstracted in such a way that network connections are formed automatically. While Java uses secure sockets, it is difficult for a client to verify that a stub is using secure sockets, potentially compromising sensitive information. Furthermore, issues such as access-control are not explicitly handled by Jini.

The Service Location Protocol (SLP) [17] is a decentralized, lightweight, scalable and extensible protocol for service (or resource) discovery within a system. SLP eliminates the need for a user to know the name of a network host, but rather, the user supplies the desired type of service and a set of attributes

which describe the service. SLP is not designed to scale to large numbers of users and its performance is questionable in unknown dynamic networks.

Universal Plug-and-Play (UPnP) is Microsoft's standard for resource discovery. Resources in the system advertise and describe themselves using the eXtensible Markup Language (XML). UPnP relies heavily on XML, HTTP, and IP and therefore can leverage known and tested communication models. For security, UPnP relies on existing World Wide Web security models such as SSL. While these methods are secure, they are computationally intensive and may not be applicable in an environment where clients and resources are computationally-starved.

The Portolano Project [16], developed at the University of Washington, is a large-scale networking infrastructure designed to support pervasive computing environments. Initial security proposals involve using IPsec-based authentication [10], but these issues have yet to be explored and remain an area of ongoing research.

6 Conclusion

This paper has experimentally verified the merits of our a resource discovery system that integrates access control by comparing it to alternative systems. The resource retrieval time is greatly reduced using this architecture, while security is not compromised. This allows our system to scale to levels that traditional resource discovery systems wishing to implement access control would be unable to efficiently reach. While the implementation and execution of this system does require additional memory in each intentional name router, sacrificing storage for time and efficiency is a worthwhile tradeoff.

Together with the proxy-based security model, this architecture meets all the goals of a secure system. It features efficient and scalable access-control for all resources while integrating with a powerful resource discovery system. It also secures all communication between proxies and the naming routers. We believe that this architecture is a flexible and generalized security infrastructure ready to support the pervasive computing trends that will surely dominate the future.

References

- [1] W. Adjie-Winoto, E. Schwartz, H. Balakrishnan, and J. Lilley. The Design and Implementation of an Intentional Naming System. *Operating Systems Review*, 34(5):186-301, December 1999.
- [2] G. Banavar, J. Beck, E. Gluzberg, J. Munson, J. Sussman, and D. Zukowski. Challenges: An Application Model for Pervasive Computing. In *Proc. ACM MOBICOM*, August 2000.
- [3] M. Burnside, D. Clarke, T. Mills, A. Maywah, S. Devadas, and R. Rivest. Proxy-based security protocols in networked mobile devices. In *Proc. ACM SAC02*, March 2002.
- [4] D. Clarke. SPKI/SDSI HTTP Server / Certificate Chain Discovery in SPKI/SDSI. Master's thesis, Massachusetts Institute of Technology, 2001.
- [5] D. Clarke, J.-E. Elien, C. Ellison, M. Fredette, A. Morcos, and R. Rivest. Certificate Chain Discovery in SPKI/SDSI. *Journal of Computer Security*, 2001. To appear.
- [6] P. Eronen. Security in the jini networking technology: A decentralized trust management approach.
- [7] P. Eronen and P. Nikander. Decentralized Jini Security. In *Proc. of the Network and Distributed System Security Symposium*, February 2001.
- [8] K. Fall and K. Varadhan. The ns manual.
- [9] Hewlett-Packard. CoolTown. See <http://cooltown.hp.com>.
- [10] S. Kent and R. Atkinson. Security architecture for the internet protocol (**rfc 2401**). November 1998.
- [11] D. L. Martin, A. J. Cheyer, and D. B. Moran. The open agent architecture: A framework for building distributed software systems. In *Applied Artificial Intelligence*, March 1999.
- [12] S. Raman. A secure framework for access-controlled resource discovery in dynamic networks. Master's thesis, Massachusetts Institute of Technology, 2002.
- [13] R. Rivest and B. Lampson. SDSI - A Simple Distributed Security Infrastructure. See <http://theory.lcs.mit.edu/~rivest/sdsi10.ps>.
- [14] Sun Microsystems Inc. Jini Network Technology. <http://www.sun.com/jini>.
- [15] UC Berkeley. The Ninja Project: Enabling Internet-scale Services from Arbitrarily Small Devices. See <http://ninja.cs.berkeley.edu>.
- [16] University of Washington. Portolano: An Expedition into Invisible Computing. See <http://portolano.cs.washington.edu>.
- [17] J. Veizades, E. Guttman, C. Perkins, and S. Kaplan. Service location protocol (**rfc 2156**). June 1997.