Proposal: Transforming Academic Computing with Public-Key Infrastructure

S.W. Smith, R. Brentrup, E. Feustel, D. Anthony, D. Nicol, L. Levine
Dartmouth College

www.cs.dartmouth.edu/˜pkilab/

Draft of November 14, 2001, 9:42

Summary

Dartmouth College proposes to develop and deploy an end-to-end public key infrastructure (PKI) for its academic user base, that overcomes the obstacles that to date have prevented PKI from taking root, and that is easily reproducible by other universities. We want to make PKI finally "happen," and to transform academic computing.

Background The technology of public-key cryptography enables two parties to engage in trusted information exchange, even if they've never met, and share no secrets a priori. Consequently, public cryptography is essential for the emerging cross-domain, electronic world, where parties need to communicate information:

- within large organizations (e.g., student X to student Y);
- across organization boundaries (e.g., student X in class C at university U_1, to professor P at university U_2);
- across time (e.g., a message from student X, club officer now, to be read by the club officer in 3 years).

An effective PKI thus enables:

- authentication (within and across institutions),
- authorization (including complex policies),
- communication of private information, and
- communication of authenticated information

across organizational boundaries, and using one basic principle.

As a consequence, standards for public-key cryptography have emerged, and numerous existing vendor and open-source products can use this technology for authentication, authorization, privacy, and digital signatures, in environments such as Web and e-mail.

Our Project What's missing is the *infrastructure* that enables large populations to use this technology. Users need certified key pairs; they need to be able to wield those key pairs from wherever they compute; and everyone needs to be able to trust that a private key is used ONLY when a user says so. That's what we want to build: "a higher-assurance PKI for higher education."

As documented in our proposal and white paper, we perceive that a number of issues have prevented the emergence of robust, extensible PKI in academic computing:

- academic users typically use many different platforms types and connect from many different locations;
- academic users (unlike, for example, military or defense users) are under no legal compulsion to follow any particular rules, and thus tend to follow the "path of least resistance"
- academia is inherently cross-domain, both internally and externally.

A effective PKI requires a user base; a user base requires applications; applications require an effective PKI.

Our project will establish a trustable, scalable PKI, and critical mass of users and applications, that works across domains, is easily manageable and reproducible, and interoperates with standard vendor and open-source products that work with public-key technology.

1 Introduction

Dartmouth College proposes that the Mellon Foundation fund a project to develop and deploy an *end-to-end*, *inter-institutional* public-key infrastructure (PKI) that:

- overcomes the limitations that hamper deployment of low-assurance PKI in academia;
- lays a foundation for developing innovative new approaches to address the security, privacy, usability, and scalability challenges that hamper the deployment of higher-assurance PKI;
- lays a foundation for exploring new directions in protection of intellectual property;
- and is easily transferable to other universities.

The goal of this project is to systematically remove the obstacles preventing effective PKI deployment for large populations. At the end of the project, we foresee:

- a large-scale, higher-assurance PKI in place at Dartmouth and other institutions;
- a critical mass of users and applications, including signed email, JSTOR/ARTSTOR, and PKI-based authentication, integrity and security and privacy for mailing list servers;
- infrastructure and application tools distributed through the Internet2 organization;
- and a steady state in which we can continue to transform academic information processes, as well as engage in further research and experimentation.

Achieving this goal requires:

- making this PKI ***usable***—so users and administrators do not need to master difficult tasks;
- making this PKI used on a ***daily basis***—so users do not need to remember tools used infrequently;
- making this PKI ***cross-institutional***—since, in the long term, PKI is the only effective means for enabling trust judgments across institutional boundaries;
- making this PKI ***secure***—since, in the long term, applications will require assurance that private key actions are authorized by the key-owning entity.

To date, widespread deployment of PKI has stalled because an effective deployment must address too many intertwining issues. As our companion white paper on transforming academic computing with PKI discusses, bringing the full vision to fruition requires understanding the information technology (IT) and process issues in deploying and maintaining PKI for a large population, and understanding where current technology falls short in flexibility, assurance, privacy, authentication, and authorization.

We believe that our project will succeed, due to our methodology. We will neither attempt to solve all the issues at once, nor pretend the issues do not exist. Rather, our work will consist of two phases, each with two concurrent tracks:

- a ***deployment track***, in which we deploy increasingly larger functionality to increasingly larger populations, with tested technology, to gain practical logistical experience with known PKI technology and methods,
- and a ***design track***, in which we develop necessary PKI technology and methods that do not yet exist.

Through the two phases of this project—and the work that will continue after—we will use the results of the design track in our next deployment track, and use questions raised from the deployment track to drive the next design track.

The applications and PKI tools we develop—as well as the lessons learned—will target cross-institutional use, and will be released throughout the project.

Timing This project is already in progress, but an expanded effort could begin as soon as funding becomes available. We anticipate that Phase 1 would consist of 12 months of work, followed by a 3-month evaluation and planning stage. Phase 2 would then consist of 9 months of work.

Section 2 will catalog more exact tasks and timeframes.

Institutions As part of our effort to have maximum impact, we have been targeting *communities* of interested universities, not just individual ones.

One part of this effort is arranging through Internet2/Educause to distribute our tools to as wide a community as possible.

Another is our leadership work in the Higher Education PKI - Technical Advisory Group (HEPKI-TAG) and Policy Advisory Group (HEPKI-PAG). When our tools are ready, we will promote them there. (Indeed, while cataloging previous university work, we discovered that there was no central contact list of who in academia was doing PKI work; we collected and organized that, as part of our efforts to build community.)

For specific testing of our work (particularly with regard to its reproducibility and its handling of cross-domain issues) we have arranged with Dr. Keith Hazelton at the University of Wisconsin, and Prof. Bennet Yee at the University of California at San Diego (UCSD) to participate in Phase 2. The Wisconsin team is experienced with PKI and Secure/Multipurpose Internet Mail Extensions (S/MIME); Prof. Yee at UCSD is experienced with secure coprocessing.

Partners More specifically, outside of Dartmouth:

- The Internet2 National Science Foundation (NSF) Middleware Initiative (NMI) program will provide a vehicle for distributing the code we develop.
- Our sister PKI Lab at Wisconsin, and UCSD, will assist us in testing the scalability and cross-domain aspects of our work.
- Dartmouth is participating in the Corporation for Research and Educational Networking (CREN) - Mellon Foundation project to catalyze PKI use for Journal Storage (JSTOR) archive access. The CREN project has garnered the interest of numerous education institutions and publishing organizations (see the Appendix).
- Dartmouth is participating in the National Institute of Health (NIH) application signing and PKI signature verification pilot using the Higher Education Bridge Certificate Authority (HEBCA) with other Internet2 member institutions and Educause members.
- Dr. John Erickson, a Digital Rights Management (DRM) expert with HP Laboratories, will also collaborate with us on the Intellectual Property(IP)-protection aspects of our project.
- We are exploring use of our framework with MIT's *Open Knowledge Initiative (OKI)* portal
- ATT/Internet2 provided the initial seed funding for the Dartmouth PKI Lab.
- Additionally, we have collaborations and equipment/code support from industrial partners, including IBM, Entrust, Baltimore, Intel, Sun/Netscape, and Microsoft (in addition to HP's contribution of Dr. Erickson's time).

Within Dartmouth:

- Prof. Denise Anthony in the Department of Sociology will be working with us on how users perceive and understand this new technology.
- The Institute for Security Technology Study at Dartmouth, under Department of Justice funding, is providing lab space and some salary and student support.
- Dartmouth Computing Services is providing servers for the PKI deployments and support of those servers.
- Dartmouth College itself (via Prof. Smith's start-up funds) has provided equipment and travel support, and is granting Prof. Smith six months of research leave in the next four years.

This Document Section 2 present specific workplans for the two phases of the project; Section 2.3 sketches the steady state that will emerge from this work. Section 3 provides more background on our workplan, and how we believe this will address crucial shortcomings in technology. Section 4 places in this implementation and deployment work in the context of other efforst by our research team. Section 6 discusses proposed staffing requirements.

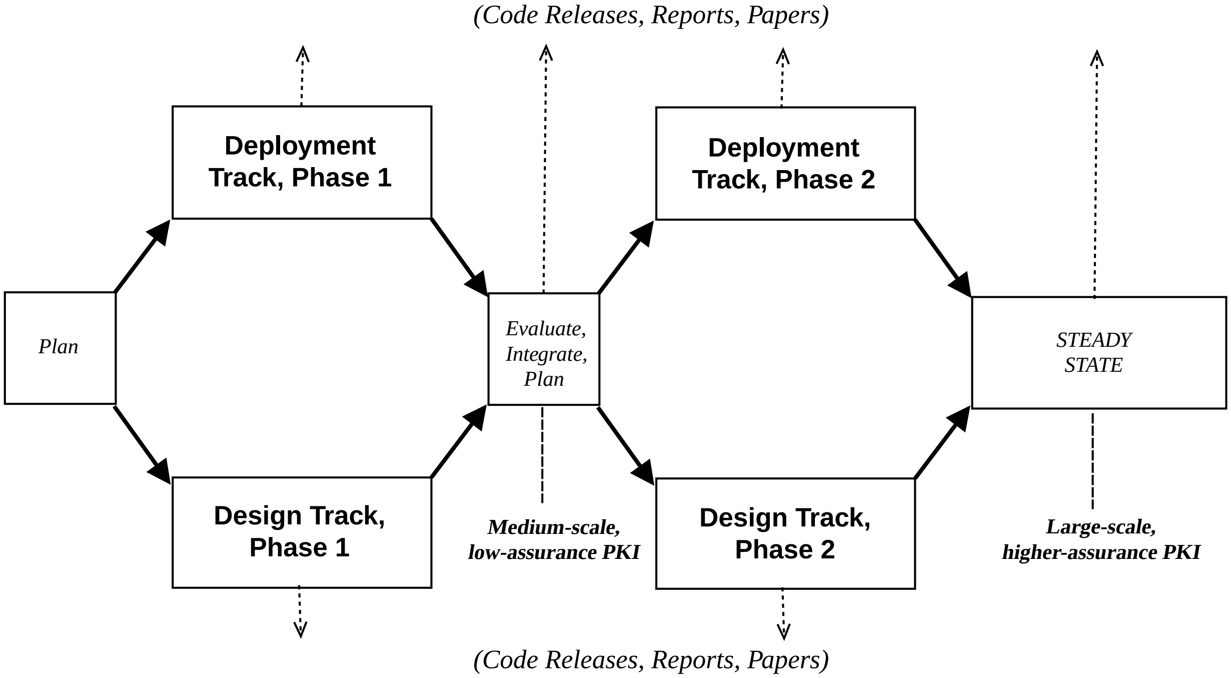
(Code Releases, Reports, Papers)
Plan
Deployment Track, Phase 1
Design Track, Phase 1
Evaluate, Integrate, Plan
Medium-scale, low-assurance PKI
Deployment Track, Phase 2
Design Track, Phase 2
STEADY STATE
Large-scale, higher-assurance PKI
(Code Releases, Reports, Papers)

Figure 1 We foresee two phases of work, interleaved with 3-month planning sessions. In each phase of our work, a *deployment track* will use current technology to investigate scalability, usability, and human logistics; concurrently, a *design track* will develop new technology to address outstanding security and privacy issues. This new technology will drive the deployment phase; the lessons learned from the current deployment will drive the next design phase. Throughout, we will generate accompanying deliverables such as code releases, reports, and scientific papers—in addition, of course, to a deployed PKI as well as PKI technology at Dartmouth and partner institutions. At the end of the process, we will achieve a steady state in which we have a higher-assurance PKI at multiple institutions, with a critical mass of users and applications. This infrastructure will function both as a daily part of campus information processes, as well as a platform for further research and experimentation.

2 Work Plan

This section catalogs the individual tasks comprising this project. Section 3.3 and Section 3.4 give additional information for the deployment and design tracks, respectively.

Percentages indicate the relevant fraction of that individual's efforts within the project.

2.1 Phase 1

In the initial deployment track, our objective is to prototype the human-oriented processes involved in an effective PKI, both to gain a better understanding, as well as to lay the foundation for subsequent phases. In the initial design track, we will focus on what we believe are the most important "missing pieces" in PKI technology.

Phase 1 Deployment Track

Initial Pilot. We will produce a short-term prototype PKI deployment to begin the establishment of a significant user population with access to a significant number of PKI enabled applications. This initial user group is needed to study operational procedures and as subjects for the human factors work.

- *Staff:* Mr. Brentrup (50%); programmer (20%); system manager (70%)

To develop this initial user group as quickly as possible, we intend to license additional commercial PKI software. (However, note that this cost is for our exploration; at the end of the projects, we plan to have free software available for universities who plan to reproduce our work.)

For this initial pilot, we plan to acquire additional client software licenses to expand Dartmouth's existing investment in its Entrust Certificate Authority, and to set up an iPlanet CMS. The iPlanet system is more cost-effective for a universal deployment and currently supports the basic applications of web authentication and secure e-mail, but weaker end-user security with its browser-stored key pairs. The Entrust software has worked well in the existing Dartmouth internal application, supports form and file signing and hardware assisted private key storage for clients. However, it requires a significant client desktop installation and is relatively expensive. Together these systems cover the range of applications, costs and complexity in PKI systems.

Signed Email Application. We will deploy signed email as one of the main applications to motivate this population. This work will also enable us to examine user interface and user expectation issues (see below), as well as to establish a mass of users to test the PKI-enhanced mailing list server and the hardened S/MIME pilots in Phase 2.

Our goal is all of these projects is universal access. This Phase 1 Deployment project will target Eudora, Outlook, Outlook Express, and Netscape. The Phase 1 Design project will target any web browser.

- *Staff:* Mr. Brentrup (30%); system manager (20%); programmer (20%)

JSTOR. As one of the pilot applications, we will deploy JSTOR access using PKI-based client authentication. We'll utilize a relatively small group of users at first, for instance, all the research library staff at Dartmouth (including student workers) because of their influence on the way faculty, staff, and students explore and use the work of others. Access would then be expanded to all potential JSTOR users using this initial group for evangelism and support.

- *Staff:* Mr. Brentrup (20%); system manager (10%)

User Study. To gain a baseline understanding of usability and user expectations, (e.g., [13] gives a computer science perspective) we will survey this user population to learn how they perceived this infrastructure and these tools.

- *Staff:* Prof. Anthony (60%); one student (100%)

From a sociological point of view, solving the technical and computational issues to ensure functionality and privacy of PKI (or any new technology) are necessary but not sufficient conditions to produce user acceptance, adoption and therefore widespread usage. Social factors (including interpersonal, institutional and economic issues) influence user perceptions of reliability and trustworthiness, as well as user decisions to begin and/or continue to use any new technology.

Recognizing this, the user study has two key components: (1) evaluating at baseline user perceptions of and behavior associated with privacy issues, e.g., user concerns about privacy protection, and user behavior regarding use of passwords, sharing of passwords, etc; (2) determining the factors that influence users to trust (i.e., accept and use) new technology, including both perceptions of computational and institutional factors, as well as social factors that contribute to assessments of reliability and trustworthiness.

(This work also ties in with the privacy study below.)

Vulnerability Analysis. The effectiveness of PKI critically depends on the fundamental axiom that an operation with an individual's private key occurs only when authorized by that individual.

To this end, we will conduct vulnerability analysis and penetration tests of the browser-based key pair approach to PKI, to better gauge its assurance (or lack thereof).

- *Staff:* Prof. Smith (20%) and two students (100%).

More specifically, how does one evaluate assertions such as "the X approach is not good enough" or "the Y approach is much better?" In the area of security, one recognized approach is to try breaking the systems: to evaluate the various attack points and strategies, and to then attempt to carry them out. (The former analysis can guide design; the latter testing can measure implementation effectiveness.)

For the particular approach of "browser-based key pairs," we will try to catalog—and then, if possible, to demonstrate—specific ways that an adversary, with varying degrees of access, can extract private keys from common browser/OS platforms. (The hypothesis is that this work will clearly demonstrate the insecurity of these approaches, and the major holes that need to be addressed. However, without this task, this hypothesis is just that—a hypothesis.)

Phase 1 Design Track

Trusted Third Party Infrastructure. As noted elsewhere, we believe a fundamental obstacle to long-term effective PKI is the inability of browser-based key pairs to support mobile users, and to provide sufficient assurance that all private key operations are authorized by the entities that own them. To address the security issues, we will use high-end *secure coprocessors* (e.g., [12]) as *trusted third-parties (TTP)* to generate, certify, and apply user key pairs. In this task, starting with open source and freeware, and our existing WebALPS hardened server project,[1] we will develop the code and tools necessary for this TTP-hardened server-side PKI. An intermediate result will be integration of S/MIME with Web-based mail; the final result will be a general-purpose PKI. (Section 3.4 provides more detail.)

[1] Our current WebALPS code is based on Apache and OpenSSL, and is slated to be open-sourced 4Q2001.

- *Subtasks*:
    - ***Design and prototype TTP-based PKI***
    - ***Build hardened S/MIME application*** by integrating this TTP PKI with web-based S/MIME
    - Extend TTP PKI to work with ***client-side SSL authentication***
- *Staff:* Prof. Smith (60%); three students (100%); integration engineer (40%)

Background on secure coprocessing. Trusted hardware plays a crucial role in the proposed design: in assurance; in privacy; in resistance to insider attack; and in the ability of arbitrary third-party sites to stand up trusted configurations.

However, we need to stress that we are not looking for a simple cryptographic accelerator or "secure place to keep a private key." Rather, we are looking for a generic *secure coprocessor* platform: a secure, authenticated place to protect *computation* in potentially hostile environments.

In particular: we're going to protect the entire server end of the SSL channel; the server-side software that authenticates the user; the usage (and any associated record-keeping) of the user's private key; and any associated cross-certification, back-up, and escrow.

Background on our platform choice. The IBM 4758 was designed to be a generic secure coprocessor platform. We (or anyone else that gets their public key certified by IBM) can sign and release software for it. The 4758 security architecture then allows any third party to

- obtain a coprocessor
- install our application
- and then have *that installation* of *that application* prove that "it's the real thing, doing the right thing," beyond the ability of anyone—including the operators of that machine, or the authors of the application—to subvert it.

The IBM 4758 pioneered this concept, and has been on the market since 1997. We chose it for several reasons:

- It was the first FIPS 140-1 Level 4 device. (To date, only IBM, Baltimore, and Thales have managed to achieve this level of validation.[2].)
- It easily permits application development by non-IBM entities (and will shortly support Linux inside the coprocessor)
- It provides a generic platform that enables third parties to securely load Dartmouth code—and enables that code to prove who it is.
- It is also generally cheaper than other alternatives, which typically offer weaker security and less (or none) authenticated, custom programming support.

Cambridge News. Recently, the BBC (and other media) has been reporting about a team at Cambridge attacking the 4758. Just to clarify (since the media tends to get things confused):

- The 4758 is a physically and logically secure computer, that can house applications. To date, the security of this platform has not been compromised.
- However, what those applications do is up to them—the 4758 gives developers an armored car, but they still have to lock the doors.

[2] See `csrc.nist.gov/cryptval/140-1/140val-all.htm`

- The Cambridge team found a hole in the CCA application that many folks use, not in the underlying platform.
- What we're proposing is developing our own application software for this platform and others. We're not using CCA, nor had we planned to.

Other Directions. Clearly, we are open to other hardware that would provide similar functionality; we are already exploring some of Baltimore's offerings.

Given the potentially fickle nature of specialized hardware, we will develop versions of our code that do not require it. (The imminent release of Linux for the 4758 will facilitate such portability.)

PKI-Enhanced List Server We will extend an open source mailing list server application to use PKI, to ensure that only authorized (e.g., properly signed) content is accepted for retransmittal. (This would overcome security and user mobility problems of current techniques which authorize based on the sender's e-mail address.) Our goal is to produce a working prototype for subsequent pilot; however, this work also raises some deeper questions, such as how we embody *intention* in the signature, how we integrate PKI with a moderator control channel, how we handle encryption, and how to tradeoff performance vs. key management if we encrypt and/or sign outbound data.

- *Staff:* programmer (40%); integration engineer (30%); Prof. Smith (15%)

Application Migration Process. As our companion white paper discusses, PKI brings new flexibility and security to existing campus information processes, as well as enables new ones not previously feasible. However, simply assigning a given application to a programmer and saying "migrate this" will not efficiently effect this transformation. In this task, we will develop a process for migration, and hopefully make manifest any PKI requirements or coding tools that are missing.

- *Staff:* programmer (20%); integration engineer (30%)

Privacy Study. Campus information processes raise issues of privacy, anonymity, and pseudonymity; coming into the mix are legal requirements, user expectations, as well as the added value enabled by new technology. To lay the foundation for subsequent work, we will investigate legal requirements and user expectations in this arena.

- *Staff:* Prof. Anthony (40%); Prof. Smith (15%); Student (100%).

This task ties in with the sociology work in the Phase 1 Deployment track. We want to determine the extent of ongoing user concerns regarding functionality, usefulness, reliability and privacy, and the impact of these concerns on usage.

This task also ties in with Prof. Smith's body of work, which can be characterized by looking at the mismatch between what people expect systems to do, and what the systems actually do. Such analysis and formal specification can identify shortcomings in current technology, as well as drive requirements for new technology.

Since much of our "higher assurance PKI for higher ed" pitch rests on notions such as assurance, trust, and privacy, it would be useful to see what community expectations really are.

Attribute Certificates. In the long run, we do not believe that mere identity certificates provide sufficient expressiveness for the authorization and access control decisions that will emerge in DRM and campus IT. In this task we will lay the foundation for exploring the use of attribute certificates for these problems, by installing and testing an attribute authority, as well as analyzing the security requirements for ARTSTOR and other DRM applications.

- *Staff:* Dr. Feustel (100%); Dr. Erickson (100%); Prof. Smith (10%).

Those who value the intellectual property they are providing will require tighter controls on who can access their materials. The first question that arises is what security model(s) will providers accept. The second question that arises is how can such model(s) be implemented in a generic manner so that as many as possible can be implemented using a parameterized technical framework.

As an example, David Wasley at the University of California Office of the president is working with over 200 content providers. It is unlikely that one access control policy fits all providers. But can one generic technical implementation given a proper set of parameters meet their needs sufficiently to induce them to provide material they would not provide otherwise?

Thus this task involves interaction with content providers, e.g., those providing JSTOR content, to determine their views on what conditions might permit the use of more recent material to a more limited audience.

Further, under what conditions would they be willing to permit delegation of a person's authorization, e.g., a professor to a student, to use material to which the first is able to use?

The second part of the task involves setting up a technical structure that can provide attributes of the user to the provider. These certified attributes may come from a variety of sources including the user themself and various authorities vouching for the user's entitlements. These attributes may come from persons delegating authority to the person responsible for requesting service.

The third part of the task involves determining a technical structure that permits the use of these attributes using the policy of the provider to determine whether the requester should be granted access based on the credentials presented.

We anticipate that the second part will be accomplished using a combination of Attribute Certificates and/or Proxy Certificates. We expect that part of this task will be to set up attribute certificate and proxy certificate generation and distribution within our PKI and to experiment with policy interpretation and enforcement, developing a standard application programming interface to be used for this function.

Ongoing Projects. The Dartmouth PKI Lab currently has other projects ongoing, including PKI integration with open-source secure shell (SSH) (to enable easy secure login throughout academia) and extending current browser technology to resist spoofing.[3] As appropriate, we plan to integrate the results of such work with our project.

2.2 Phase 2 Work Plan

In the second phase, we will deploy a wider, higher-assurance PKI with tools and applications built in the first phase, and we will use the lessons learned in Phase 1 to continue to work on the missing pieces.

Phase 2 Deployment Track

Larger Pilot. Building on our experience with the initial pilot and the results of our TTP development, we plan to deploy the higher-assurance TTP technology from Phase 1, both for a large part of our campus population, and for populations at partner institutions.

- *Staff:* Mr. Brentrup (60%); system manager (20%); programmer (20%); one or more students (50%) and faculty/staff (10%); University of Wisconsin colleagues, TBD; UCSD colleagues, TBD.

[3] See `www.cs.dartmouth.edu/~pkilab/demos/spoofing/`

Roll-out of PKI-Enhanced List Server and Hardened S/MIME. As part of this larger roll-out, we will deploy the PKI-enhanced mailing list server and Hardened S/MIME applications developed in Phase 1.

- *Staff:* Mr. Brentrup (40%); system manager (20%); integration engineer (20%); student (50%)

Ongoing User Study. We will continue to examine the usability and user perception of our technology.

- *Staff:* Prof. Anthony (100%); student (100%)

Test Application Migration Process. We will validate the application migration process developed in Phase 1, by using it to migrate selected legacy applications.

- *Staff:* programmer (40%)

Test Attribute Deployment. To better understand the attribute certificate approach to authorization, we will deploy AA on a small-scale with two simple DRM applications.

- *Staff:* Dr. Feustel (50%); system manager (20%); programmer (20%)

Phase 2 Design Track

Escrow. We will extend our armored vault techniques [4] to balance privacy with functionality for key escrow in academia; and prototype solutions with the TTP framework. This work will be based on user data (e.g., “how often do users forget keys?”) and privacy expectations learned in Phase 1.

- *Staff:* Prof. Smith (30%); one student (100%); integration engineer (20%)

Authentication and Delegation. We will extend our work on why students share passwords and ID cards now [2], and develop experimental approaches to building a PKI that incorporates this delegation while maintaining security. We will also use the computational power and flexibility of the TTP/coprocessor approach to develop experimental approaches to user authentication (e.g., [1]) that resist such “password sharing”.

- *Staff:* Prof. Smith (30%); two students (100%)

Our preliminary study here at Dartmouth showed that students regularly shared authenticators such as passwords and ID cards, but primarily for *delegation*: student A needed student B to carry out some specific task (such as register for room draw, or purchase a meal) in the name of student A.

Many of the benefits of PKI would be reduced, should students carry out this same behavior with the passphrases or tokens that authenticate or arm usage of their private keys.

Potential solutions here include:

- building an authentication infrastructure that makes it easy to delegate authority for such actions (so students do that instead of sharing authenticators)

- deploying authentication mechanisms (such as Berkeley's visual authentication project, where users authenticate by selecting, from a sequence of randomly chosen "genetic art" images, which belong to "their set") that make it difficult or impossible to share passwords.

However, bringing such solutions into a real infrastructure requires the ability to implement and easily tune such mechanisms, and then test how they work with real users in real situations. Our TTP approach gives us a centralized and (compared to smart cards or dongles) powerful platform on which to do this work.

Another avenue we can explore is limiting the damage should a user authenticator be exposed via a compromised browser—since the private key and its use is guarded within a trusted box running software *we* develop, we can easily extend this software to address various scenarios, such as one-use passwords and hash-chaining within signatures.

Privacy for the List Server We will experiment with using TTP and PKI techniques to bring privacy and anonymity to the mailing list server: e.g., using a TTP to validate a sender's signature, but then stripping the signature and resigning it with the mailing list server's private key, would enable authenticated but anonymous content, even against some types of insider attack.

- *Staff:* Prof. Smith (20%); one student (100%); programmer (20%)

Advanced DRM Applications. We will prototype advanced applications that attempt to use attribute certificates or proxy certificates to enforce IP policy for ARTSTOR and other projects; this may build on our ongoing secure browser work.

- *Staff:* Dr. Feustel (50%); Dr. Erickson (100%); Prof. Smith (20%); one student (100%)

This presence of DRM in a PKI proposal follows from two points:

- a critical mass of applications are needed for PKI to take root
- in the long run, PKI is the only way to express authorization and trust across organizational boundaries (as well as time and space)

DRM is an arena where both of these issues come into play—as well as being relevant to the central role of academia: collection and dissemination of information. That is why we have an emphasis on using DRM to motivate PKI, as well as looking at PKI to enable new approaches to DRM.

2.3 The Future

Availability

The Dartmouth team wants our work to make an impact in the real world.

To this end, we will ensure that:

- Any software we produce will be open source and available for public use.
- If any patentable ideas emerge during this work, we will either place them in the public domain, or pursue patent protection but ensure that the Mellon Foundation will be fully licensed to use them.

Furthermore, we have opened discussions with IBM to determine if any patents pending from Dr. Smith's prior work there (e.g., the early WebALPS idea) may impact this project, and, if so to arrange full licensing for educational use. (Precedent suggests there will be no obstacles here.)

Sustainability

Following the completion of this work, the various PKI implementations in production will be sustained by Dartmouth College Computing Services/Technical Services. Computing Services recognizes the need for the services described in this proposal. Dr. Larry Levine, Dartmouth's Director of Computing ("CIO"), has endorsed this effort by his direct participation. Dartmouth College would like to be a center for work in the area of PKI, serving all of higher education.

Furthermore, this project will create a critical mass of PKI users, applications, and infrastructure that will enable ongoing research that addresses the open problems in this area and improves the efficiency of open source implementations of solutions to them. The direct participation of Prof. David Nicol, Chair of the Computer Science Department and a Program Director at the Institute for Security Technology Studies, and the support already provided by these institutions, testifies to their endorsement of this work.

Higher Education Management

Effective PKI tools for higher education will greatly facilitate the transaction of academic and administrative processes. PKI speaks to many of higher education's needs for the secure control, access, and usage tracking of digital content. As is widely acknowledged, the traditional information paradigm which was exclusively print-based now includes the exponentially growing world of digital information. In addition to the direct challenges of creating and adapting digital information for various scholarly and administrative uses, issues of secure access are increasingly vital in related realms of information management. These include:

- managing communication services (such as email),
- "published" intellectual property,
- cost-effective sharing of information,
- presenting proprietary and confidential information securely and variably (especially for instance via course and learning management systems, and "portals"),
- transacting necessary organizational processes on-line,
- interacting with the federal government,
- and tracking utilization.

The work proposed here will help all of higher education to concertedly determine and implement PKI services.

3 Strategy

In this section, we discuss the evolution of our strategy and workplan.

3.1 Current Obstacles

The combination of a number of issues has limited the rate at which PKI has been adopted at universities, and in society at large. These issues range from basic logistics to fundamental research; but success in all these issues is necessary to go beyond mere demonstrations of PKI, and instead achieve wide-scale deployment and usefulness.

- **The Fundamental Axiom.**

  The effectiveness of PKI critically depends on the fundamental axiom that: ***an operation with an individual's private key occurs only when authorized by that individual.***[4] A PKI deployment that does not provide this fundamental axiom is essentially pointless.

  However, the inability to assure that this axiom holds even against moderate adversaries continually hampers deployment of PKI and non-trivial applications based on it. Providing this assurance—while working within usability and economic constraints—raises many challenges, including:

  - deciding where the private keys for individual users (and other entities) should live, and how they should be safeguarded;
  - evaluating the resiliency of laptop and desktop keystores to attack from remote adversaries;
  - providing secure private key services for mobile and remote access users;
  - balancing security, flexibility, and cost of tokens containing and protecting keys;
  - overcoming the strong tendency of individuals currently to share passwords.

- **Applications.** A primary issue is creating applications sufficiently compelling to motivate campus populations to use PKI and for campus IT to invest the time and effort for implementing and maintaining it. This issue is a common stumbling block:

  - applications require PKI,
  - but PKI requires applications.

  PKI is the only secure access paradigm with sufficient potential to fulfill the emerging variety of needs for IP and digital rights management (DRM) access control—including balancing privacy of access with accountability (for billing, logistics, etc.)

  Our experience in campus IT support has clearly demonstrated that users will not master tools unless they use them daily. For example, in our payroll authentication pilot, users could not remember passphrases they only used once a month! Even a low-assurance PKI deployment will not thrive, if JSTOR is the only application.

- **Logistics.** Discussion of PKI often begins with cryptography, and overlooks the logistical issues required to deploy and maintain PKI for a large population of individuals. For example, campus IT staff must provide an adequate infrastructure for checking identity when enrolling users, and for verifying identity after enrollment.

- **Economics.** Several vendors currently provide technology for many PKI components: for example, Lightweight Directory Access Protocol (LDAP)-enabled directory servers, certificate authorities, and certificate management systems. However, vendor solutions are typically too expensive and too inflexible for the open, decentralized model of inter-institutional academic computing. Open-source solutions are also available, but require more integration and extensions.

[4] The term "nonrepudiation" is sometimes used to denote exclusively the legal aspects of this issue, and sometimes to denote the broader topic. We avoid the term here because of this ambiguity.

- **Heterogeneity.** An effective campus PKI must accommodate a wide variety of machine types, operating systems, and client software. Already, individual users expect to interact with the campus information infrastructure from a wide variety of platforms. On campus there are stationary desktops, typically with a single user; there are mobile laptops and smaller computers that are typically associated with a single user, and others that come from shared pools; and there are public access terminals. Off-campus, there are computers used by staff, faculty and students at their homes or while traveling which can remotely access a university's network, and still fruitfully be part of the university PKI.
- **User Mobility.** An effective campus PKI must also permit any given user to interact from a variety of platforms. Already, students use lab machines, public-access machines in coffee houses, shared laptops, and personal machines; traveling staff use machines at conferences and other universities.
- **Privacy.** Many aspects of PKI inevitably lead to balancing user *privacy* against other social goods. For one example, usability and legal concerns often requires key escrow—but that potentially exposes private keys to abuse. For another example, using identity certificates for authentication reveals user identities; but using other certificate schemes that preserve privacy can hamper accountability.
- **Trust.** The goal of PKI is enabling effective ***trust judgments***: e.g., "the individual at the browser is authorized to see this journal article" or "the server at the other end of the wire is authorized to receive my private health information." At their root, these judgments stem from human decision and human policy. Consequently, an effective PKI must provide client tools to enable human users to make effective trust judgments, and server tools to enforce and enable human administrators to express trust policies.

3.2 Our Strategy

Designing and deploying an effective PKI can thus lead an organization into a "Catch-22" situation:

- A near-term deployment does not allow time to address the security and privacy issues in the PKI or applications based on it;
- However, a design that addresses such issues requires a near-term deployment of infrastructure and applications to establish sufficient critical mass to allow non-trivial treatment of these issues.

Our two-track, multi-phase approach, working with partners at other institutions, resolves this dilemma. In each phase, our deployment track will explore the scalability, usability, and human logistic issues, and our design track will explore new technology to address security and privacy issues left unresolved by the currently deployed technology Each subsequent phase will integrate the results of the previous phase's prototypes.

This approach gives us the best of both worlds:

- Starting with a near-term deployment shortens the time to a short term prototype deployment, lowers risk, and allows us to learn from mistakes before deploying the PKI on a large scale.
- However, also starting with longer-focus research efforts means that we will not be saddled with a dead-end, insecure system, and we have the ability to treat both the Phase 1 deployment and design efforts as prototypes.

3.3 Deployment Track

In our first phase, we will produce a short-term prototype deployment with browser-stored key pairs to begin the establishment of a user population and to evaluate approaches to the ***logistical*** issues:

- establishing CA/Registration Authority (RA) roles, policies, and certificate validity periods;

- determining registration procedures and systems;
- evaluating certificate and key renewal procedure;
- generating usability studies (e.g., why do students share IDs and passwords? how often, and under what conditions, do they forget passphrases?);
- testing operational and compatibility problems;
- and testing inter-institutional use of directories for certificate lookup.

We'll focus on just a few applications first, and on supporting a monotonic security model (where certificates expire, but aren't revoked.) We'll utilize a relatively small group of users first—for instance, the research library staff at Dartmouth, including student workers, because of their influence on the way faculty, staff, and students explore and use the work of others. In subsequent efforts, we will broaden this pilot to include larger campus populations and more applications, as well as to incorporate the results of our design work.

Our initial application will provide a client population PKI-based authorization to JSTOR. Locally, we will also establish some low-risk data access applications (e.g., balances on the Dartmouth debit card), as well as inter-institutional shared Web site access, and inter-campus signed S/MIME e-mail.

Preparation. To prepare for this work, we are currently setting up our research staff to use PKI on a regular basis; identifying partners; setting up PKI-based authorization for our private PKI Lab site; and experimenting with S/MIME and PGP email.

We're also integrating our legacy Entrust system with LDAP-enabled directories while getting a small group set up with Entrust tools and Eudora, and preparing to enroll with the CREN CA.

Tasks. To set up a basic PKI for our initial pilot, we foresee the following tasks:

- ***Obtain certificate authority status with CREN.*** This requires us to establish security policies and practices that are appropriate to Dartmouth and satisfy CREN. Much of the work here involves determining where the Dartmouth private key will be stored, how it will be kept safe, and how it will be used to sign public key certificates for users.
- ***Establish a registration process and a registration authority.*** Since we want to make certain that this can be scaled to include the whole University, this will require substantial work with other organizations to develop a process that they can administer.
- ***Determine an appropriate period of certificate validity***. This period must meet the needs of JSTOR as well as Dartmouth. It is essential that the cost of reissue be minimized, especially if the period is short, or if the cost of externally supplied certificates is high.
- ***Determine user requirements and expectations for key pair privacy and accessibility.*** In this phase, we will use current browser-based key pair storage, with no escrow. However, in order to plan our subsequent design phases, we need to understand what user requirements and expectations exist in this area.
- ***Establish a process for registering public key certificates in our LDAP-enabled directory.*** This process must assure that the PKC is entered registered to the correct user and that it is not counterfeit.
- ***Establish a backup and recovery process for our LDAP-enabled directory and key databases.***
- ***Determine a basic Application Program Interface (API)*** for extraction of information from certificates for use in policy-based authentication and authorization.

- ***Determine compatibility.*** Determine which browsers/mailers and/or clients for applications will be supported and which will not.
- ***Developing training material*** and a schedule of training sessions for the Certification Authority, Registration Authority, Application Developers, and Users. When possible we will develop web-based training exercises.

We will perform vulnerability assessments and penetration tests of this infrastructure.

In the second phase, we plan to deploy the higher-assurance TTP technology from the Phase 1 design track, both for a large part of our campus population, and for populations at partner institutions. We will work on additional campus applications.

3.4 Design Track

In our design track, we will focus on designing and prototyping an infrastructure that provides a foundation for addressing the mobility, security, privacy, and non-repudiation issues that hamper effective PKI. For example, browser-stored key pairs create many obstacles for security, user mobility, and heterogeneity; entity key pairs alone limit the flexibility of PKI to support advanced rights management authorization policy.

Security To address the security issues, we will use high-end *secure coprocessors* (e.g., [12]) as *trusted third-parties (TTP)* to generate, certify, and apply user key pairs. The advantages of such an approach are manifold:

- By moving the private key away from the client platform, we reduce the risk of exposure from permeable desktop machines, and also increase client mobility.
- By moving the keys into high-assurance hardware, we reduce the risk of exposure to adversaries—including the server operator.
- By also moving sensitive computation—not just the private keys—into trusted hardware, we gain additional end-to-end security.
- By centralizing the high-assurance hardware, we get economies of scale over a user-token approach—as well as the benefits of vastly improved physical security.
- By using a high-end programmable platform, we gain flexibility for exploring creative approaches to user authentication, time-stamping and hash-chaining of key operations, key escrow, and key revocation.

We'd rather think of this hardware as a "super-dongle" than as a proxy. Portable dongles and user tokens limit user mobility, offer weaker security[5] than a high-end coprocessor, and offer a much more inflexible, limited programming environment to support experimentation.

Preparation. Our trusted third-party approach minimizes system development risk by building on our past work. Prof. Smith invented and coded much of the underlying security technology in the IBM 4758 coprocessor, [12]; Our team has already developed an infrastructure to use coprocessor-based TTPs for *Armored Vaults* [4], for WebALPS Web servers protected against insider attack [5, 6, 10], and have started designing a system to use this technology for integrating S/MIME with Web-based mail, without exposing user keys to remote machines, or the server operator [7].

[5] Our team has unique expertise in this matter; S.W. Smith led the software and formal modeling effort that earned the first-ever FIPS 140-1 Level 4 validation [11].

TTP Work. We want to have TTP prototypes that work with campus users, and work with remote applications (not just JSTOR) that expect client-side SSL-authentication. In the initial design work, we plan the following explicit tasks:

- ***Port WebALPS server code*** into Linux/4758 environment, so users with legacy browsers can open secure channels into the trusted third party. (An open-source Linux release for the 4758 is imminent [3].)
- ***Design and implement outbound authentication support*** for the Linux/4758 environment, so that users can tell what application is in there. (Prof. Smith designed and co-implemented this for the current CPQ/4758 environment [9].)
- ***Set up a code-signing facility*** (getting a public key certified by IBM) so that we can make application software available for secure installation by anyone.
- ***Set up a CA that uses outbound authentication*** to verify WebALPS installations, and then issue them SSL server certificates.
- ***Extend our current client security work*** [14, 15] so that users can securely verify identify and attributes of servers, despite attempts by malicious servers to spoof.
- ***Extend our hardened S/MIME*** work [7] to develop a WebALPS application that generates key pairs for users, certifies them, and uses them in an S/MIME compatible e-mail system.
- ***Design and prototype a scheme*** where the above TTP scheme can efficiently work with third-party servers, such as JSTOR, that expect SSL client-side authentication.

Authorization. To address the authorization and access issues, we will explore the use of attribute and proxy certificates to enable the expression of short-term rights independent of the longer lasting identity certificates to which they are tied. We will also lay the groundwork for a more extensive evaluation of the ways in which policies for use of digital information can be enforced. The JSTOR community can serve as one of the communities through which such policies and the means of their enforcement can be pursued. We are investigating "ready made" products that provide attribute certificates and APIs providing access to the attributes and values in these certificates. We are preparing to develop these capabilities if they are not commercially available.

We expect that proper protection of intellectual property will require policy-based management of that property. Further, in order to pass credentials with confidentiality and integrity, attribute and proxy certificates or their XML counterparts—both of which require key pairs—will be required. Because authorizations usually have a shorter duration than credentials issued that confirm identity, it is likely that special forms of credentials will be required for digital rights management. In addition, it is likely that negotiation between client and server will be required to determine what credentials are to be exposed and at what time. In the event that consumable rights are employed, it may be that transactions that can be "rolled back" will be needed.

In our initial work we will develop/purchase an attribute certificate authority and deploy it in small scale. We will determine whether it can be used with current browsers, possibly by means of plugins using the Netscape Personal Security Manager (PSM). Additional work to modify the PSM data schema may be required. We will also coordinate with the JSTOR project to try to understand how attribute certificates might meet their current and future needs.

We will work with content providers to determine what kinds of authorization policies they would like to enforce. We will serve as technology interpreters in explaining what is doable and what the doable costs. We will experiment with a variety of authorization policies and enforcement technologies that are centered around public/private key pairs. The goal of the advanced stage is to permit a single set of mechanisms to serve the largest possible collections of application and data providers as well as their users. By using these mechanisms, application and data providers may well be willing to permit the use of applications and data either for free or at lower cost because they will know that they can provide fine grained access control for their intellectual property.

Subsequent Research. In subsequent work, we plan to use the TTP base to experiment with escrow and authentication techniques, as discussed in the workplan.

We also plan to experiment with the use of attribute certificates to address security requirements in DRM applications such JSTOR and ARTSTOR. We also need to understand many things about attribute certificates and their use in protecting intellectual property and in conveying entitlements to this property.

Additional issues—to be explored during the project, and beyond—include:

- Integrating PKI with freeware SSH and campus machine infrastructure, to enable users to securely sign in (`ssh`) to academic machines anywhere
- Integrating PKI with more complex Web information services such as those requiring community of interest access control.
- Building on our current Marianas work [8] to examine how a campus-wide distributed network of TTPs can support typical distributions of users.
- Examining approaches for "instant revocation" by inhibiting TTP from operating for that key pair.
- Extending our previous work on using TTPs for auctions, into *credential negotiation.*
- Integrating PKI into an effective solution for "drive-by networking" (that is, unauthorized connection to an institution's wireless network)

To ensure usability of our technologies, we will explore application development tools that facilitate using attribute certificates.

3.5 Impact

The immediate outcome of this project will be a preliminary PKI deployment at a small number of institutions:

- that is used for JSTOR, and likely other content and information services,
- that provides a working foundation for addressing the usability, mobility, and security issues that have hampered PKI to date,
- that provides usage data, to identify subsequent areas for tuning and development,
- and that generates supported software deliverables to ensure that our results, with the potential for this project to provide continuing support for institutions that may adopt these tools.

Our dual emphasis on both the logistics of deployment, and the security of the resulting system is critical. ***Deployment is pointless if the system is easily subvertable; but higher security is pointless, if the system is never deployed or is used infrequently.***

Deployment The deployment of a PKI also supports a growing need—relevant at Dartmouth as well as other institutions—to increase the number of transactions conducted in a secure manner, on both the Internet and internal institutional networks. The operational objectives here are to improve service by reducing the time for action completion and reduce costs by reducing the need for additional administrative personnel. A campus-wide supported security infrastructure should in turn simplify application development and its costs.

Improved network-based administrative services will meet customer expectations for the availability of self-service applications and allow location independence as faculty and students are frequently off campus. Secondarily, a PKI should reduce the “time to market” for newly purchased directory-enabled applications by providing standard compatible interfaces to authentication, authorization, accounting, and auditing services.

In many ways, these types of services are competitive advantages for attracting faculty, staff, and students. They may also be necessary to work with systems envisioned by the Federal Government for Student Loans and Grants and Contracts applications and processing.

Project Goals Two related sets of goals will be achieved by this project. The first is to develop a process and software product that enables the deployment of a medium assurance PKI that a university can follow which is easy to understand and can be implemented by a typical university IT operation. This implementation is to be available at a cost affordable by a typical university and permit at least the three following applications: S/MIME signed/encrypted mail with Outlook, Outlook Express, Netscape 4.x, and Eudora with Tumbleweed plugin; Access control based on Apache using Access Control Lists and client certificate values; JSTOR utilization. When possible, the process will admit the use of commercial products in addition to open source software for those universities requiring a higher level of assurance or having other considerations. An API that can be used by new/converted applications will be provided so that a university can modify its own applications or create new ones. The process and software will permit inter-university applications as well as intra-university applications.

Second, the successful completion of this project will integrate the pilot PKI deployment and the advanced security design to deploy a higher-assurance PKI on campus-wide basis. This will provide a basis and some initial answers for examining the many open questions that need to be answered before an effective universal PKI takes hold. The applications and code developed will both be tested in production at Dartmouth and offered to our collaborators to assist creating the critical mass needed to have the technology used on a daily basis.

The results of this project will enable deployments capable of authenticating system users, assisting in determining authorization rights, encrypting data transmissions, securing documents and interoperating with external institutions. We will research various commercial product offerings select and install the most promising, as well as prototype and build missing pieces. We will detail the steps and resources needed scale to it to the entire university communities and demonstrate how a university can utilize the Higher Education - Federal Government trust bridge with the PKI. We will work with partner institutions as they take similar steps and with others interested in developing capabilities based on PKI in exploring mutual interests.

Working out compatible data definitions for operation between institutions, as well as defining the needed policies and procedures and local operational details are major challenges of this project.

Deliverables This is a list of some of the tangible outputs we currently envision for the project, from both the deployment and design tracks. They are described in various places in the proposal and are gathered here in summary:

- Distribution of Dartmouth developed PKI tools and applications
  - Trusted Third Party (TTP) infrastructure
  - WebALPS application to store key pairs for S/MIME users
  - Integration of PKI with freeware SSH
  - Secure mailing list server implementation
  - WebALPS code for CPQ/4758, and Linux/4758
- Multiple Institutions using PKI access control to JSTOR
- JSTOR Access Support Web site
  - Sample implementation roadmap
  - Report to guide other PKI deployments

  - Recommendations on:
    - PKI Products and Compatibility
    - Certificate Authority Policies and Procedures
    - Registration Policies and Procedures
    - Key delivery and protection
    - Certificate and Key renewal issues and procedures
    - Directory considerations
    - End user studies
    - Operational issues, staffing and costs
  - Training materials for end users/ help desk consultants/ developers
  - Considerations for additional Service Providers

- Documentation on PKI Lab web site
  - Studies on attribute certificates, multiple end-user key pairs, key escrow and revocation issues
  - Study of Root models, hierarchy or bridge
  - Resources for application authors
  - documentation, tools, papers, pointers to other useful PKI sites

- Certificate Authorities
  - Code-signing server for our IBM 4758 academic PKI applications
  - CA (under CREN) to authenticate and issue SSL server certificates for our and other institution's IBM 4758 academic PKI applications

- Progress and Status Reports to Mellon

4 Research Context of Work

This proposal focuses on making effective PKI happen in academia. This work requires some research to fill holes; however, we have been careful to restrict that work to using the results of research that is already completed, or is already underway but sufficiently low-risk. Research is just one component in a much larger implementation and engineering portfolio presented here.

It is true that the research arm of our team have long-standing interests in secure coprocessing and PKI, and our research proposals reflect this. In contrast to those longer-term exploratory proposals, the present proposal focuses on implementation and deployment (although we hope, in the long run, for synergy between this practical work and our more abstract efforts).

The Dartmouth team is composed of both IT and Research personnel; successful completion of this project would further all of our goals: those of Mellon, of Dartmouth IT, and of Dartmouth research.

5 Relation to the WebISO Project

The WebISO project uses Web cookies to address the problem of how to convey authorization from a central server, through a client browser, to a remote application that expects this means of authorization. In comparison to our project, WebISO is an orthogonal technology, addressing a much smaller problem: conveyance of local authorization. WebISO itself does not address how users authenticate, nor how to extend across multiple domains (see section 4 of the 8/2/2001 requirements draft). WebISO also does not address the other problems that PKI solves: privacy and integrity of communication between peers (e.g., signed and encrypted email, or signed documents, such as homework submission, payroll authorization, grants and contracts, etc.)

In contrast, our PKI project lays the foundation for secure electronic business and education within and across institutions.

In subsequent discussions with the WebISO team, we have considered the implications of the two proposals and believe that they are complementary and even synergistic.

First, the proposals differ in the technical areas they're addressing. WebISO enables establishment of authenticated web sessions within an institution; Dartmouth's focuses on the additional data privacy, integrity, and trust applications enabled by PKI. WebISO focuses on code that can be shipped immediately; Dartmouth's work initially will build the foundation for shippable tools within the year. WebISO focuses primarily on securing applications within an organization (leaving cross-institutional operations to other systems); Dartmouth's work focuses on the unique ability of PKI to communicate trust across a broad set of boundaries, including institutional boundaries. WebISO is layered on top of existing security infrastructure; Dartmouth is trying to maximize the robustness of the foundation for the next-generation infrastructure.

Furthermore, the proposals are synergistic; each will benefit from the success of the other. WebISO is a first step for institutions with no coherent security infrastructure to transition to a stronger system and does not specify architecture or security practices for the weblogin server; Dartmouth's distributed Trustable Third Party technology can plug those gaps. Dartmouth is moving user key pairs out of the browser, and will require a way to authorize user access to their key pairs through their browser; WebISO's technology (with extremely short-lived cookies) can help handle this problem.

Finally, WebISO uses PKI to provide security for its server-to-server messages. Dartmouth's PKI work will make it possible to easily and securely craft and process such signed and encrypted statements about trust, across multiple institutions. Consequently, Dartmouth's work should help enable transition of WebISO to higher-assurance, cross-domain settings.

6 Staff and Resources

6.1 Staffing

To staff the project we expect to be able to acquire a greater portion of the time of Prof. Smith, Dr. Feustel and Mr. Brentrup for the effort funded in part by the project. Prof. Smith will lead the research track and oversee the work of the graduate and undergraduate students; his primary focus will be TTP-based computing, and security and vulnerability analysis. Mr. Brentrup will lead the deployment track and serve as the project manager. Dr. Feustel will lead the investigation of authorization and delegation policies and technologies including the use of attribute certificates or other methods of conveying credentials from client to server using public key cryptography.

We note that all three members of this core group have experience in industrial software and tool production.

Additionally, as we have discussed, we are working with Prof. Anthony in Dartmouth's Department of Sociology to help design surveys and experiments relating to alternative technologies used in PKI and monitor the implementation of human processes required for the implementation of PKI. This collaboration will enable us to understand the human issues behind the transition from paper-based human processes to digital processes as automation of applications proceed—at Dartmouth, at other institutions, and in society as a whole.

The core technical group will also include numerous undergraduate and graduate students under the supervision of Prof. Smith and Dr. Feustel, and will be expanded with the employment of a System Manager, a System Programmer and an Integration Engineer to support the efforts of both tracks. In addition, Dartmouth has staff assigned to related projects—such as directory, library and administrative systems—who may be able to assist.

Recruiting for these new project positions would begin once funding could be determined. The positions could be phased in a sequence. The system manager would be the first position to fill, followed by the system programmer and then the integration engineer. While these types of positions have been difficult to fill at times, there are currently many more applicants for similar work available. While for continuity it would be preferable to engage dedicated staff for the duration of the project, it may be possible to fill these roles in a number of creative ways, including multiple part time or temporary arrangements in an effort to get underway quickly. Short descriptions for these positions are included below.

6.2 Job Descriptions

Project Manager Plan and track Mellon/CREN and JSTOR projects, supervise project personnel; identify application areas; work with application project managers to link into Middleware systems; oversee system architecture; participate in the evaluation and selection of hardware and software components for PKI.

System Manager Install, operate, update and maintain PKI server software (Entrust, Unicert, CDSA, iPlanet) and underlying server hardware; also Middleware directory systems; support production systems and testbed systems for research; assist application programmers with integration, testing and production rollout; maintain documentation and resource links on PKI Lab web site; work with IT security officer to maintain security of PKI.

Programmer Develop connector components to complete PKI systems; including remote PKI enrollment system for students; carry out mailing list server modifications develop shared code for application systems, such as code for authorization checks (e.g., an individual's group membership); extend open source systems.

Integration Engineer Complete and integrate student projects; clean-up and package results of internal projects for distribution; contact for deployment at other institutions, assists with remote installs and troubleshooting; maintain demos and documentation sets on PKI Lab web site.

References

[1] R. Dhamija and A. Perrig. “Deja Vu: A User Study Using Images for Authentication.” *9th USENIX Security Symposium.* August 2000.

[2] E. Etu and J. McIsaac. *Bringing PKI to Dartmouth.* Class Project, CS88, Dartmouth College. June 2001. (In preparation for submission.)

[3] “IBM Research Demonstrates Linux Running on Secure Cryptographic Coprocessor.”Press release, August 28, 2001.

[4] A. Iliev. *An Armored Data Vault.* Senior Honors Thesis, Department of Computer Science, Dartmouth College, June 1, 2001. Dartmouth CS Technical Report TR2001-400.
`www.cs.dartmouth.edu/˜pkilab/papers/alex.ps`

[5] S. Jiang. *WebALPS Implementation and Performance Study.* (Master’s Thesis). Computer Science Technical Report TR2001-399, Dartmouth College. June 2001.
`www.cs.dartmouth.edu/˜pkilab/papers/shan.ps`

[6] S. Jiang, S.W. Smith, K. Minami. “Securing Web Servers against Insider Attack.” *ACM/ACSA Annual Computer Security Applications Conference.* December 2001. (To appear.)
`www.cs.dartmouth.edu/˜pkilab/papers/jsm.ps`

[7] E. Knop. *Secure Public-Key Services for Web-Based Mail.* Senior Honors Thesis, Department of Computer Science, Dartmouth College, August 2001.
`www.cs.dartmouth.edu/˜pkilab/papers/evan.ps`

[8] D.M. Nicol, S.W. Smith. C. Hawblitzel, E. Feustel, J. Marchesini, B.S. Yee. “Survivable Trust for Critical Infrastructure.” *Internet2 Collaborative Computing in Higher Education: Peer-to-Peer and Beyond.* 2001. (To appear.)

[9] S.W. Smith *Outbound Authentication for Programmable Secure Coprocessors.* Computer Science Technical Report TR2001-401, Dartmouth College. March 2001
`www.cs.dartmouth.edu/˜pkilab/papers/oa.pdf`.

[10] S.W. Smith. “WebALPS: A Survey of E-Commerce Privacy and Security Applications.” *ACM SIGecom Exchanges.* Volume 2.3, September 2001.
`www.cs.dartmouth.edu/˜pkilab/papers/acm.pdf`

[11] S.W. Smith, R. Perez, S.H. Weingart, V. Austel. “Validating a High-Performance, Programmable Secure Coprocessor.” *22nd National Information Systems Security Conference.* October 1999.
`www.cs.dartmouth.edu/˜sws/papers/nfips.pdf`

[12] S.W. Smith, S.H. Weingart. “Building a High-Performance, Programmable Secure Coprocessor.” *Computer Networks (Special Issue on Computer Network Security.)* 31: 831-860. April 1999.

[13] A. Whitten, J.D. Tygar. “Why Johnny Can’t Encrypt: A Usability Evaluation of PGP 5.0.” *USENIX Security*, 1999.

[14] E.Z. Ye. *Securing Browsers.* Master’s Thesis, Department of Computer Science, Dartmouth College. In progress.

[15] Y. Yuan, E. Ye, S.W. Smith. *Web Spoofing 2001.* Computer Science Technical Report TR2001-409, Dartmouth College.
`www.cs.dartmouth.edu/˜pkilab/papers/spoofing.ps`

Appendix A

CREN/JSTOR Participants

- Columbia University
- Dartmouth College
- Georgia Institute of Technology
- MIT
- Oberlin College
- Princeton University
- University of Chicago Library
- University of Illinois Urbana-Champaign
- University of California
- University of Minnesota
- University of Texas
- JSTOR
- Stanford University HighWire Press
- OCLC Online Computer Library Center

Dartmouth is also participating in the Internet2 HEPKI projects in which at least twelveinstitutions (see below) have been regularly participating in a wide variety of PKI volunteer effort activities[6] (Note that there is substantial overlap with the CREN project group.) Through the Mellon project Dartmouth could focus this group's interest in S/MIME e-mail, enlisting them as participants in pilot deployments.

HEPKI Regular Participants

- Cornell University
- CREN
- Dartmouth College
- Georgetown University
- Georgia State University
- MIT
- University of Alabama
- University of Calfornia
- University of Massachusetts
- University of Michigan

[6] See `middleware.internet2.edu/` and `www.educause.edu/hepki/`.

- University of Virginia
- University of Washington
- University of Wisconsin

Appendix B: The Dartmouth Team

At Dartmouth College, our Computing Services group, the Computer Science Department, and the Institute for Security Technology Studies are pursuing a research agenda that identifies and addresses the security and privacy issues in the current and emerging computational infrastructure. In October 2000, this group was selected to host one of two national PKI Labs to further research in PKI issues. This award and past work on support of Academic computing as detailed subsequently make a compelling case that Dartmouth is the correct institution to support for further development of academic applications of PKI.

Organizations

Computing Services. Our Computing Services team has already deployed a limited PKI to selected staff in support of particular vertical applications.

Developing a campus wide PKI is in progress and being pursued in conjunction with the Internet2 Middleware Early Adopters project.[7] The Early Adopters project goals are to develop best practices and documentation for other educational institutions to follow when deploying middleware.

Computing Services also operates a directory infrastructure for security and rights management for a large number of users, spanning organizations as diverse as other universities, medical centers, and a local non-profit ISP. It has deployed a a Kerberos infrastructure covering Dartmouth College and Dartmouth Hitchcock Medical Center (DHMC) associated Medical Centers. It has implemented a distributed authentication and authorization protocol (IDAP) that permits secure sharing of information resources among multiple campuses. Several institutions jointly license database resources that are mounted on servers at Dartmouth College. Many current administrative and library applications requiring authorization are supported at Dartmouth using various combinations of Dartmouth Name Directory, Kerberos and IDAP. It is recognized that the current ad-hoc nature of determining user attributes limits the development of additional applications. We are also finding that commercial groupware applications like calendaring require directory and/or PKI middleware to function.

Dartmouth's Computing Services has a long history of developing innovative, networked services and effectively deploying them across the entire campus. Examples of networked services developed at Dartmouth include the Potlatch file server (a circa 1985 high-performance AppleShare and FTP server), Dartmouth Name Directory (a circa 1986 campus-wide white pages and authentication directory), BlitzMail (a circa 1987 institutional e-mail system), Fetch (a circa 1989 Macintosh FTP client), the Dartmouth College Information System, DCIS (a circa 1991 campus-wide client-server information ”dashboard”), and InterMapper (a circa 1996 network management application). All of the above services are in active campus-wide use in 2000 and their developers are available as technical resources.[8]

Department of Computer Science. Computer science has a long history at Dartmouth College, from early groundbreaking research in the 1960's through the bachelor and doctoral programs of today. In the early 1960s, Dartmouth became one of the first institutions to make computers easily available to every student and faculty member. With its development of the BASIC programming language, Dartmouth completely transformed academic computing internationally.

In 1979, Dartmouth created the undergraduate major in Computer Science; In 1993, the computer science faculty moved to the new Sudikoff Laboratory for Computer Science, and in 1994, formally split from the Department of Mathematics.

Dartmouth's teaching mission provides a sophisticated, questioning user base to push the limits of what we build and provide requirements for new directions. Our faculty are known world-wide for their excellence in research

[7] See `http://www.internet2.edu/middleware/earlyadopters/`

[8] See `http://www.dartmouth.edu/pages/softdev/`

and are recognized as some of Dartmouth's best teachers (and Dartmouth was recently ranked #1 in teaching among universities nationwide). Our graduate students are the recipients of numerous prestigious graduate fellowships and their work is published in the major computer-science publications. Our undergraduate students are among the best in the world, and many reach beyond the classroom to join faculty in research projects or to join computer corporations in developing next-generation computer software.

ISTS. Dartmouth is home to the new Institute for Security Technology Studies (ISTS). ISTS was recently established to serve as a national center for research in security technologies related to cyber-security and information infrastructure protection. Funded at $15M/year, the Institute's technical agenda is led by researchers from Dartmouth's Department of Computer Science, Thayer School of Engineering, and School of Medicine. The Institute also encompasses partners from leading universities and laboratories. ISTS contributes a unique resident group focused on security technology, as well as existing collaborations with over a dozen other expert groups.

The ISTS is a research center funded through the Department of Justice, within the National Institute of Justice (NIJ). By needs it must focus on aspects of security that most directly fall within the scope of NIJ, which are law enforcement, particularly at the state and local level. While ISTS enjoys a large budget, only a very small portion of that budget can be allocated for PKI-related work. Furthermore, that PKI work must be directly and obviously tied to the NIJ mission. It is not yet obvious to NIJ that PKI does this; ISTS is working to educate them on this point, by focusing its PKI-related effort on showing how specific applications in law enforcement are served by a PKI, e.g., construction of a distributed crime investigation team, an internal investigation team, coordination of sensitive data among diverse state and local organizations. Early applications of PKI in this context will have to be very simple, because we must rely upon relatively unsophisticated system administration infrastructure in the client agencies.

By contrast, the academic applications envisioned for the Mellon Grant are substantively different, and promise to involve significantly more complexity. The level of system administration expertise one can expect at an academic institution is higher, which simply makes it possible to explore richer avenues.

Key Personnel

The personnel leading this effort represent a cross-section of Computing Services, Computer Science, and ISTS.

Sean Smith (Assistant Professor, Computer Science) has been working in information security—attacks and defenses, for industry and government—for over a decade. In graduate school, he worked with the US Postal Inspection Service on postal meter fraud; as a post-doc and staff member at Los Alamos National Laboratory, he performed security reviews, designs, analyses, and briefings for clients including USPS, SSA, HHS, INS, and DOE; at IBM T.J. Watson Research Center, he designed the security architecture for (and helped code and test) the IBM 4758 secure coprocessor, and then led the formal modeling and verification work that earned it the world's first FIPS 140-1 Level 4 security validation. In the field of computer security alone, Dr. Smith has published over a dozen refereed papers; given over two dozen invited talks, and has over a dozen pending patents; his security architecture is used in thousands of financial, e-commerce, and rights managements installations world-wide. Dr. Smith was educated at *Princeton* (A.B., Mathematics, 1987) and *CMU* (MS, Ph.D., Computer Science) and is a member of ACM, USENIX, Phi Beta Kappa, and Sigma Xi.

Robert Brentrup is the Associate Director of Technical Services for Dartmouth College Computing Services, working on networked computer applications including e-mail, directories, authentication and authorization systems. Previously he was the project director for the Dartmouth College Information System (DCIS) and managed Library Information Systems. Prior to that, Robert worked at Lotus Development Corp., where he was a principal engineer involved in the Lotus Jazz, Notes and Improv products, Spartacus Computers and Raytheon Co. He is the author of a number of professional papers and was a member of the Northeastern University faculty. Robert holds a B.S. in Electrical Engineering from *Michigan Technological University* (1977) and a M.S. in Computer Engineering from *Boston University* (1981).

Edward A. Feustel is a Research Associate at Dartmouth's Institute for Security Technology Studies. Ed's research focuses on security issues in distributed computing applications and infrastructure. He previously was employed in the Computer and Software Engineering Division of the Institute for Defense Analyses, Alexandria, Virginia where he was a Research Staff Member with interests in Security and Distributed Systems and Applications. Prior to IDA Alexandria, he worked at Prime Computer as a Principal Technical Consultant, Rice University as a tenured Associate Professor of Electrical Engineering and Computer Science, IDA Princeton as a Systems Programmer, Lawrence Livermore Laboratory as a Sabatical Researcher, and California Institute of Technology as a Research Fellow. He has represented Prime and IDA as technical liaison to the Object Management Group (OMG), to the Open Software Foundation (OSF), and X/Open. He is a graduate of *Princeton* (MA 1966; Ph.D. 1967 in Electrical Engineering) and *MIT* (BSEE and MSEE, 1964). His 1971 IEEE Transactions of Computers Paper on the Advantages of Tagged Architectures set the stage for network exchange of tagged data such as CORBA and application exchange of tagged data such as XML.

Denise L. Anthony received her Ph.D. in Sociology from the University of Connecticut in 1997. During 1997-1999 she was a Robert Wood Johnson Foundation Scholar in Health Policy Research at the University of Michigan. Since July of 1999, she is Assistant Professor in the Department of Sociology, and Adjunct Assistant Professor in the Department of Community and Family Medicine, at Dartmouth College. Dr. Anthony's research interests include collective action problems, organizational theory, economic sociology and the sociology of health care. She explores how incentive and control structures, including economic mechanisms as well as social relationships and informal social norms, affect behavior, organizational outcomes, and organizational and institutional change. For example, she has studied the effects of social relationships in diverse settings, from the changing risk-norms among networks of injection drug users, to the process of trust formation and cooperation among members of micro-credit borrowing groups, to the changing nature of physician referral relationships under managed care. In addition, Dr. Anthony is currently studying the adoption and effects of new rules and practices by managed care firms on patients, providers and markets.

David Nicol is Professor and Chair of the Department of Computer Science at Dartmouth College, as well as the Technical Coordinator for the Cyber-security program at ISTS. He has long worked on issues in systems, and distributed and parallel computing; this work is balanced by consulting work with IBM, NASA, AT&T, and Sandia National Laboratory. He publishes prolifically, and serves as the Editor-in-Chief of ACM Transactions on Modeling and Computer Simulation. He holds a B.A. in Mathematics from *Carleton College* (1979), and a Ph.D. in Computer Science from the *University of Virginia* (1985).

Larry Levine has been Director of Computing at Dartmouth College since 1991. He oversees an enterprise-wide IT environment, including academic, administrative and network IT functions. Prior to his current position, Larry held a variety of management positions within Dartmouth and as a programmer, research consultant, and manager at Indiana University's Computing Services (1979-1984). He holds a Ph.D. from *Indiana University*, Bloomington, where his major focus was on research methodology and statistics, and a B.S. from *S.U.N.Y. Stony Brook.*