

3rd Annual PKI R&D Workshop
April 12-14, 2004, NIST, Gaithersburg MD

Online Proceedings

The **official proceedings** are published, and can be ordered, as [NISTIR 7122](#)

This online version of the proceedings reflects the original program, and includes everything in the official proceedings (papers and summaries) as well as those presentation materials that have been provided by the presenters.

Workshop Summary

Ben Chinowsky, *Internet2*

Monday, April 12 2004

12:00 - 12:15

Opening Remarks

Program Committee

Dr. Susan Zevin, *Director, Information Technology Laboratory, NIST*

12:15 - 1:15

KEYNOTE ADDRESS

Non-Intrusive Cross-Domain Identity Management

Stefan Brands, *Credentica & McGill University*

2:00 - 3:30

Role Sharing in Password-Enabled PKI

Xunhua Wang, *James Madison University*

Greenpass: Decentralized, PKI-based Authorization for Wireless LANs

Nicholas C. Goffee, *Dartmouth College*

X.509 Proxy Certificates for Dynamic Delegation

Von Welch, *National Center for Supercomputing Applications, University of Illinois*

4:00 - 5:30

Panel: Controlled and Dynamic Delegation of Rights

Frank Siebenlist, *Argonne National Laboratory*

Carl Ellison, *Microsoft*

Kent E. Seamons, *Brigham Young University; Internet Security Research Lab*

Ravi Pandya, *Microsoft*

Von Welch, *National Center for Supercomputing Applications, University of Illinois*

8:00 - 9:00

Work in Progress Session

Ben Chinowsky, *Internet2*

Tuesday, April 13 2004

9:00 - 10:00

KEYNOTE ADDRESS

How to build a PKI that works

Peter Gutmann, *University of Auckland*

10:00 - 10:30

Experiences of establishing trust in a distributed system operated by mutually distrusting parties

David Chadwick, *University of Salford*

11:00 - 12:15

Panel: NIH-EDUCAUSE PKI Interoperability Project: Phase Three

Peter Alterman, *National Institutes of Health*

Russ Weiser, *Betrusted*

Scott Rea, *Identrus*

Deb Blanchard, *Digital Signature Trust*

12:15 - 12:45

Trusted Archiving

Santosh Chokhani, *Orion Security Solutions*

Carl Wallace, *Orion Security Solutions*

12:45 - 2:00: Lunch - NIST Cafeteria

2:00 - 3:00

Panel: Document Signatures

Randy Sabett, *Cooley Godward LLP*

John Landwehr, *Adobe*

Ron Usher, *Juricert*

3:00 - 3:30

PKI: Ten Years Later

Carlisle Adams, *University of Ottawa*

4:00 - 4:30

An Examination of Asserted PKI Issues and Proposed Alternatives

John Linn, *RSA Laboratories*

4:30 - 5:30

Panel: Which PKI Approach for Which Application Domain?

Peter Alterman, *National Institutes of Health*

Carl Ellison, *Microsoft*

Russ Weiser, *Betrusted*

8:00 - 9:00

Birds of a Feather sessions

Wednesday, April 14 2004

9:00 - 10:00

KEYNOTE ADDRESS

A New and Improved Unified Field Theory of Trust

Ken Klingenstein, *University of Colorado; Internet2*

10:00 - 10:30

Private Revocation Test using Oblivious Membership Evaluation Protocol

Hiroaki Kikuchi, *Tokai University*

11:00 - 11:30

"Dynamic Bridge" Concept Paper

Ken Stillson, *Mitretek Systems*

11:30 - 12:45

Panel: Approaches to Certificate Path Discovery

Peter Hesse, *Gemini Security Solutions*

Steve Hanna, *Sun Microsystems*

Matt Cooper, *Orion Security Solutions*

Ken Stillson, *Mitretek Systems*

12:45 - 2:00: Lunch - NIST Cafeteria

2:00 - 2:30

Johnson & Johnson Use of Public Key Technology

Rich Guida, *Johnson & Johnson*

2:30 - 3:00

Identifying and Overcoming Obstacles to PKI Deployment and Usage

Jean Pawluk, *Inovant*

3:30 - 4:30

Panel: The PKI Action Plan: Will it make a difference?

Steve Hanna, *Sun Microsystems*

Sean Smith, *Dartmouth College*

John Linn, *RSA Laboratories*

Lieutenant Commander Thomas Winnenberg, *U.S.N., DISA*

Tim Polk, *NIST*

Jean Pawluk, *Inovant*

4:30 - 5:00

Wrapup Session



3rd Annual PKI R&D Workshop Summary

[Ben Chinowsky](#), *Internet2*

The workshop announcement listed the goals of this gathering as:

1. **Explore the current state of public key technology** in different domains including web services, grid technologies, authentication systems et al. in academia & research, government and industry.
2. **Share & discuss lessons learned and scenarios** from vendors and practitioners on current deployments.
3. Provide a forum for leading security researchers to **explore the issues relevant to the PKI space** in areas of security management, identity, trust, policy, authentication and authorization.

This summary groups workshop sessions according to which of these goals was their primary concern, although many sessions addressed more than one.

Surveying Deployments

Dr. Susan Zevin, Director of the Information Technology Laboratory at NIST, opened the meeting by noting some Federal PKI highlights. The Federal Bridge Certification Authority now connects six Federal agencies. The Department of Defense now requires contractors to obtain PKI credentials for email and authentication to DoD web sites. Several countries and international associations, including Canada, Australia, and NATO, are negotiating to connect to the Federal PKI. NIST is a global leader in smartcards and biometrics and their integration with PKI.

A panel discussion with Peter Alterman, Deb Blanchard, Russ Weiser, and Scott Rea discussed the **NIH-EDUCAUSE PKI Interoperability Project: Phase Three**. This project has been largely driven by the Government Paperwork Elimination Act; in order for virtual paperwork not to become just as much of a hassle as the physical paperwork it replaces, reducing the number of certificates each person needs ("reusability") is essential. While this is still at the technology-demonstration stage — a production implementation has additional, expensive, datacenter requirements — various agencies including GSA and HHS are adopting elements for production use. This uptake in the federal environment is what this seed project is all about. The panelists' [project report](#) describes progress to date in great detail.

The use of **Document Signatures** in land-ownership transactions in Canada was also the subject of a panel discussion. Attorney and former crypto engineer Randy Sabett compared physical and digital signatures, and in particular explored the issue of digital signatures being held to higher standards than physical ones. John Landwehr, from Adobe, described how Acrobat supports signatures from both the author and the user of a form, in order to guard against spoofing and data modification respectively; there has been strong customer demand for this. The centerpiece of this panel was [Ron Usher](#)'s description of the application of the tools described by Landwehr to a real-world situation that raised many of the legal issues described by Sabett: moving the documentation of Canadian land-ownership transactions to an electronic format. Forgery of paper documents has been a big problem in the Canadian land-tenure system; this and the need for greater efficiency were the principal drivers of the move to secure electronic documentation. Usher described his philosophy as PKE, with the E standing for Enough: "usually what we really need is public-key cryptography," with infrastructure to be added only as needed. Usher's company, Juricert, was launched by the Law Society of Canada to implement this approach. Lawyers, not landowners, are the ones who sign the documents in the Canadian system, so it's only they who need certificates. On the other hand, Usher observed that lawyers tend to be very conservative about process. One key to user acceptance of the switch to electronic transactions is to make the form look as much like the paper version as possible. This is a main reason for choosing Acrobat (though a TIFF image is the permanent legal record). The new system provides an "astounding improvement" in transaction time. The government had been re-keying information keyed and printed by lawyers; this system eliminates the keystroking — a big win for the cash-strapped government. The benefits have prevailed over the lawyers' conservatism: the new system has handled \$400 million (Canadian) in offers and ownership transfers in the few weeks it has been in operation.

[Rich Guida](#) offered an overview of **Johnson & Johnson's Use of Public Key Technology**. The J&J PKI is enterprise-directory-centric — a certificate subscriber *must* be in the enterprise directory (which is an internally LDAP-accessible AD forest). Guida stressed the importance of providing proper training for helpdesk personnel and providing enough helpdesk resources. J&J produced a one-page document distilling what users need to know to use the PKI — what tokens are for, where you use them, what to do when asked for a passphrase, etc. — and found that users often wouldn't read even this, but would instead call the helpdesk for even the most basic questions. On the other hand, J&J was able to do most configuration and credential preparation independently of the users. Guida also noted that while it has taken significant effort to get users to treat their USB tokens as a must-have item like their car keys or J&J badge, "the beauty of using the token is starting to catch on." Users particularly appreciate having a single passphrase that doesn't have to be complex or be changed every 90 days. USB tokens were chosen over smartcards only because of the ubiquity of USB ports; Guida expects a move to multifunction smartcards (e.g., used for building access also) over time. Standardization on 2048-bit keys will help drive the transition.

[David Chadwick](#) related **Experiences of establishing trust in a distributed system operated by mutually distrusting parties**. The mutually distrusting parties in question are national governments involved in a worldwide effort to monitor production of environmental contaminants capable of doing harm across international borders. Currently about 100 of 300 monitoring sites are sending signed messages to a data collection center. Every message must be signed by a threshold number of the mutually distrusting parties; this m-out-of-n principle is used wherever possible. Chadwick noted that human factors have been a major focus in both deployment and operation.

There were also two presentations relating experiences using PKI for the specific tasks of delegation and archiving.

[Von Welch](#) reviewed the use of **X.509 Proxy Certificates for Dynamic Delegation**. Proxy certificates were first prototyped in 1998 and were standardized in PKIX earlier this year; an RFC is imminent. Proxy certificates are part of the Globus toolkit and are now widely used in scientific testbeds in many countries. There are three authorization models: identity-based authorization (i.e., impersonation), restricted delegation of rights, and attribute assertions without delegation; most implementation experience has been with the first of these. The users seem pleased; their main complaint is that certificates exist as files on the local machine.

In the **Trusted Archiving** session, [Santosh Chokhani and Carl Wallace](#) described a proof-of-concept trusted archive that they built for the US Marine Corps. The approach taken was refreshed timestamps, with RFC 3161 rules used to verify that the correct data was stored. Chokhani called the group's attention to LTANS, an IETF working group formed for trusted archive standards.

Drawing Lessons

Two sessions were devoted primarily to this goal.

Peter Gutmann keynoted on **How to build a PKI that works**. After presenting an entertaining catalogue of PKI disasters, Gutmann offered a list of six "Grand Challenges" for PKI, along with proposed approaches to meeting those challenges.

1. Challenge: key lookup. Response: "the Web is the Public File." In its simplest form, this would mean putting a certificate on your home page and letting people find it with Google; while he's not promoting this, Gutmann noted it would still be better than anything currently available. His more serious proposal is "http glue + anything you want"; pretty much any database now supports the Web, many with surprisingly little effort. See <http://www.ietf.cnri.reston.va.us/internet-drafts/draft-ietf-pkix-certstore-http-07.txt>.
2. Challenge: enrollment. Response: make it transparent. Gutmann quoted Bo Leuf: "the vast majority of users detest anything they must configure and tweak." The norm when trying to get a certificate issued is to be subjected to pages and pages of hassle; there is a persistent myth that this is inherent in the process of certificate issuance. By

contrast Gutmann cited the ISP model: you call the ISP with a credit card, they give you a username and password, you use them, DHCP does the rest. We need to remember that our PKI-enabled applications only have to be as secure as the best non-PKI alternative. More on this "plug-and-play" approach to PKI is in <http://www.cs.auckland.ac.nz/~pgut001/pubs/usenix03.pdf>.

3. Challenge: validity checking. Response: Gutmann outlined an approach based on querying hashes submitted by users; this puts the work on the client.
4. Challenge: user identification. Response: Distinguished Names "provide the illusion of order" but create chaos. Gutmann used a variety of examples of this to argue for treating Distinguished Names as meaningless bit strings, and using binary compare for name comparisons.
5. Challenge: no quality control. "Some of the stuff out there is truly shocking." Again Gutmann provided a rich variety of examples. Response: Create "brands" and test procedures to become brand-certified (e.g., S/MIME testing under RSADSI); against these brands, test the basics only (lookup, verification and basicConstraints/keyUsage enforcement); make sure that users know that while software certified to the brand will work, software not so certified could do anything.
6. Challenge: implementer/user apathy. E.g., never updating CRLs, but checking against them anyway in order to formally meet requirements. Response: "Make the right way the only way to do it."

[Gutmann's slides](#) for the workshop (124 of them) develop his proposed approach in detail; he also provides crypto libraries to support it (see <http://www.cs.auckland.ac.nz/~pgut001/cryptlib/>).

The other big "lessons learned" compilation was [Carlisle Adams'](#) presentation on **PKI: Ten Years Later**. Adams dates PKI from 1993 and the beginnings of X.509 dissemination. Three big lessons from those ten years are:

- As Gutmann detailed, PKI is hard to implement.
- User issues are key.
- The principal lesson of the many deployment issues is the need for many choices in the various PKI technology areas. In each of the six principal areas of PKI technology — authority, issuance, termination, anchor, private key, validation — the last ten years have increased the number of choices from one to several. The key now is to use this large toolkit for real-world deployments.

There were three particularly interesting exchanges in the Q&A portion of Adams' session:

- Adams' reference to Ellison and Schneier's "10 Risks of PKI" as the best-known compilation of criticisms of PKI (along with Gutmann's, which is more deployment-oriented) prompted Ellison to point out that he himself is now a critic of that paper. (See <http://world.std.com/~cme/html/spki.html> for links to this paper, and to the CACM Inside Risks columns derived from it, which Ellison considers to be better written.) Ellison noted that he and Schneier were directing their fire primarily at the marketing literature around PKI, not at PKI technology itself; he recommended his paper from PKI02 (<http://www.cs.dartmouth.edu/~pki02/Ellison/paper.pdf>) as a substitute. Schneier, however, still pushes the "10 Risks" paper.
- In response to David Chadwick's observation that DNS demonstrates the viability of a global namespace, Ellison predicted that political forces will never allow a global namespace to happen again. Owning the world's namespace gives you tremendous power; it happened the first time because nobody noticed until it was already established, and because it was created by technical people who weren't out for political power.
- Eric Norman suggested that the one thing that's remained constant over the last ten years is the keypair. Adams replied that not even that has remained constant — once people thought everyone just needed one keypair or maybe two (for signing and encryption); now there's a general acknowledgment that everyone will need multiple keypairs.

Identifying Tasks

The bulk of the sessions at PKI04 were devoted to identifying and prioritizing tasks needed to move PKI forward. The two main themes that emerged from the previously described sessions — 1) human factors and 2) letting practical needs drive technology choices rather than vice versa — were dominant here as well.

Six sessions addressed directions for specific technical areas.

In the **Controlled and Dynamic Delegation of Rights** panel, participants put forward various tools for addressing this problem. Moderator [Frank Siebenlist](#) presented on Grid needs for delegation of rights; he believes that industry will face similar requirements in two or three years. Carl Ellison argued that when rights are delegated it is vital that the act of delegation be performed by the authority on the rights being delegated, rather than by the party that happens to control some authorization database. More generally, Ellison stressed that the user is part of the security protocol; Ellison's work on procedures designed to take proper account of this fact ("ceremonies") is documented in http://www.upnp.org/download/standardizeddcp/UPnPSecurityCeremonies_1_0secure.pdf. Ravi Pandya presented XrML 2.x as a general-purpose policy language, not the narrow DRM language it's often seen as (XrML 1.2 was much more limited). [Kent Seamons](#) presented TrustBuilder (<http://isrl.cs.byu.edu/projects.html>), an architecture for automated trust negotiation based on gradual disclosure of credentials. Seamons noted that this is a growing field; William Winsborough is another key person working in this area.

In the discussion, Steve Hanna asked why there had been no presentation on XACML; Frank Siebenlist, who's on the XACML TC, noted that XACML has no delegation capability, though there are plans to add this. Carl Ellison related his experiences with SPKI/SDSI to his current involvement with XrML: lack of industrial-strength toolkits and marketing are the main reasons SPKI hasn't deployed; this in turn is due to SPKI's lack of CAs precluding anyone from making money from it. But, XrML has all the power of SPKI/SDSI and more, and it's backed by Microsoft. Pandya added that the core of XrML is pretty much final, and that toolkits are in the works. Microsoft is committed to getting organizations like Globus to take it up and work it to its full broad potential. Information on the IPR status of XrML is at <http://www.xrml.org/faq.asp>.

[Ken Stillson](#) of Mitretek presented a "**Dynamic Bridge**" **Concept Paper**. Stillson observed that the path-discovery process scales very poorly and is brittle: path discovery has no sense of direction, and taking a wrong turn can lead to a wild goose chase. "Namespaces aren't organized in a way that facilitates a routing protocol." The Dynamic Bridge provides a means of consolidating paths so that intermediate nodes no longer make you have to guess. There is substantial overlap between these ideas and work on shortening certificate chains done by Radia Perlman at Sun. Mahantesh Halappanavar noted that he and his co-authors have also published work along similar lines. Mitretek owns the patents on the Dynamic Bridge concept, but has no intent to assert patent protection. They are looking to start a discussion on possibilities for implementation; contact stillson@mitretek.org if you are interested.

Stillson's talk was followed by a panel discussion on **Approaches to Certificate Path Discovery**. [Peter Hesse](#) reviewed the basic PKI structures that path discovery must deal with, describing them as all meshes, just of different shapes. Path building has not yet been addressed by IETF standards, but an informational Internet-Draft (I-D) is in the works. Steve Hanna explored analogies for path building. Is it graph theory? Only if you download all the certificates in the world. Is it navigation? Sort of. Really it's like building a deck — going out and getting things, then repeatedly running back for things you forgot, is most of the work. So, work with what you've got, keep scraps, collect tools ahead of time, and work carefully. The common theological issue of the right direction in which to build paths needs to be answered accordingly: "it depends." Meeting in the middle is also an option, particularly appropriate for bridged topologies. Hanna suggested that more research is needed: test different path-discovery modules with different topologies, and try to find the best algorithms for particular sets of circumstances. This would make a great master's thesis and could generate dozens of papers. Matt Cooper summarized the approaches he's taken in writing three pathbuilding modules, and shared test results quantifying the usefulness of various simplifications such as pruning and disallowing loops through CAs. He also stressed the importance of doing as much validation as you can in the process of doing discovery. Ken Stillson

stressed that in addition to the tasks of path discovery and path validation there is also the task of object location — as there is no global directory, even if you know the Distinguished Name (DN), you don't necessarily know how to get the certificate, so you end up having to implement a bunch of different access methods.

Hesse then moderated a discussion:

What is the goal when discovering paths? The consensus here was that (as Hanna put it) "any path is a good path." Cooper observed that it's likely that the first path you find is the intended path even if it's not valid, so that path should be reported to the user. It's also important to be able to specify a timeout: e.g. users only want it to take a few seconds for email, and a search that takes more than five minutes is very unlikely to succeed.

Is path discovery best done on the client or on the server? There appears to be a consensus that the answer here is the same as the answer to the forward vs. backward issue — "it depends" — though Stillson pointed out that audit requirements may dictate doing path discovery on the server.

What are your recommendations for PKI architects?

- Hanna: Send "a bag of certs" to the end entity via S/MIME or SSL; use name constraints in cross certificates; avoid high fan-out/fan-in.
- Stillson: Take advantage of the documents coming out of NIST. These include recommendations drawn from trying to get the bridge to work, in particular on certificate profiles, directory structure, and path discovery requirements.
- Cooper: Use name constraints; put data in the directory where it belongs.
- Hesse: Make sure your keyIDs match; use the authorityInformationAccess field.

Who has the obligation to do path discovery? The only consensus on this appears to be that it is an important unresolved question. Stillson noted a related question: Who's liable if a valid path tells me to do something I shouldn't?

What can be learned from PGP? Hesse observed that PGP doesn't really have a discovery mechanism; the user needs to know somebody it trusts, then build a local copy of the PKI that it cares about. On the other hand, Stillson cited the trust scoring system in PGP as having relevance. Neal McBurnett pointed the group to statistics on the PGP web of trust and links to path-building services at <http://bcn.boulder.co.us/~neal/pgpstat/>.

Steve Hanna wrapped up the path-discovery session by asking all with sample PKI topologies to send them to him (shanna@funk.com) for testing. Anyone interested in further research on path discovery and validation should also contact him.

[Nicholas Goffee](#) presented **Greenpass: Decentralized, PKI-based Authorization for Wireless LANs**. This project is driven by guests wanting access to Dartmouth's wireless network. Greenpass uses a SPKI/SDSI authorization certificate to bind authorizations to a public key; the delegation process makes use of a "visual fingerprint" assigned to a guest and verified by the delegator before signing the certificate. The certificate chain gets stored as a cookie on the guest's machine so the guest can reauthorize without repeating the introduction process. A pilot deployment is in the works.

[Xunhua Wang](#) presented a method for **Role Sharing in Password-Enabled PKI**. Roles are preferred to individuals as the subject of security because they are more permanent and because security policies are concerned with roles, not individuals. The principal advantage of the proposed approach is its lightweightsness: users need passwords only, not smartcards or segments of the private key.

[Hiroaki Kikuchi](#) outlined a **Private Revocation Test using Oblivious Membership Evaluation Protocol**. In the course of certificate status checking, OCSP servers learn the relationship between the certificate holder and certificate

checker. There is a privacy issue here; the proposal outlines an "oblivious membership test" to address this.

Another six sessions were specifically devoted to identifying key issues and next steps for PKI as a whole.

[Stefan Brands](#) outlined a comprehensive heterodox approach to making use of public-key cryptography: **Non-Intrusive Cross-Domain Identity Management**. In Brands' view, the Achilles heel of X.509 is its fundamental incompatibility with privacy: public keys are "strongly authenticated 'super-SSNs'". Brands pointed out the shortcomings of various proposed solutions to the privacy problem within the X.509 framework: pseudonyms and roles, attribute certificates, per-domain CAs and certificates, and federated identity management. Instead, "new authN primitives" are required. Brands' alternative, called Digital Credentials, is based on twenty years of research by dozens of academics, starting with David Chaum's work in the 1980s. The features of Digital Credentials include "sliders" for privacy and security, selective disclosure/hiding of attributes, unlinkability, and a variety of choices along the identity-pseudonymity-anonymity spectrum. Digital Credentials are patent-protected, but Brands stressed that this is only so that he can secure the investments needed to drive real-world deployments. Brands is willing to make the technology available where doing so does not conflict with this goal; contact him if you have ideas for collaboration. Brands' ideas are developed at length in his book, *Rethinking Public Key Infrastructures: Building in Privacy*.

[John Linn](#) of RSA offered **An Examination of Asserted PKI Issues and Proposed Alternatives**. Linn's proposed alternatives are more along the lines of new ways of using X.509: Identity-Based Encryption and related approaches; online trusted third parties; distributed computation; alternative approaches to validation (hash trees in particular); key servers; and privacy protection via pseudonyms and attribute certs. Linn also noted that "you can't have full success until you've had partial success first," and that choices such as hierarchical vs. nonhierarchical PKIs — once matters of ideological controversy — are now matters of pragmatic adaptation to circumstances.

In a panel discussion on the question **Which PKI Approach for Which Application Domain?**, Peter Alterman, Carl Ellison, and [Russ Weiser](#) explored some of the specifics of this latter point. The theme of PKI not being a one-size-fits-all technology, but rather a technology that needs to be custom-tailored to a huge variety of real-world situations, has become steadily more prominent over the last couple of years, and the contrast between this session and the "Dueling Theologies" session at PKI02 (<http://www.cs.dartmouth.edu/~pki02/theologies.shtml>) illustrates this nicely. Ellison stated his continuing belief in the importance of local naming — not so much to avoid name collisions, which can be addressed by domain component (dc) naming, but in order to provide a means of "the relying party knowing who this is." The relying party needs to be able to use a name it assigns — a name it can remember — for a trusted entity. Ellison claims that SPKI/SDSI and XrML can do everything needed here; X.509 might work if the environment is constrained accordingly. Rich Guida (the other dueling theologian from PKI02) observed that there's increasing recognition that if you want to join a namespace, you have to choose between changing your naming or not joining; conflicts should be resolved at join-time. The problem is that you still have to have a way of knowing who others really are, what they call themselves; relying on entities to attest to the identity of others is inescapable.

Guida suggested that doctors, for instance, would never bother to assign a local name for every patient with whom they'd need to securely exchange information. This led into a discussion of PKI in medical scenarios more generally. Peter Gutmann observed that doctors confronted with PKI usually just sign in at the start of the day and let everyone else use the machine. Doctors rightly don't want anything in the way of their work; you have to design around the fact that they see any kind of security as an impediment. PDAs that transmit certificates to the network, and short-range RFIDs, were suggested as approaches to security in emergency rooms and similar settings. Guida suggested that PKI will be used a lot more in the context of medical research and clinical trials, where there isn't the "get the certificate vs. restart the patient's heart" problem, but where there is a strong need to ensure data authenticity, integrity and confidentiality. Another possible application is finding out if a suspected drug-of-abuse-seeking patient has been to other clinics. Ellison pointed out that this use case requires an aggregator, but — contrary to common perception — doesn't

require X.509, or any other particular variety of PKI. No global name for the patient is needed; what matters is that the aggregator have one, and only one, key for each patient.

[Ken Klingenstein](#) keynoted on **A New and Improved Unified Field Theory of Trust**. Klingenstein identified three spheres in which individuals require trust: personal, work, and transactions where extremely high assurance is needed (often transactions involving the government). For each of these, there is a type of trust relationship which is usually appropriate: peer-to-peer, federations, and hierarchical PKI, respectively. Virtual organizations cut across these boundaries and thereby represent an additional challenge. Klingenstein described P2P trust as "a bedrock of human existence;" expressing it in electronic form is therefore necessary. It's also hard, although PGP, webs of trust, and X.509 proxy certificates have made some progress. Federations are getting deployed; Merck has a large and noteworthy deployment. Klingenstein noted that the federation structure for InCommon will be per-nation, as attitudes and practices for security are nation- and culture-specific. InCommon is hoping to set up a virtuous circle between the use of trust and the strengthening of trust. Klingenstein also offered an overview of recent developments and ongoing projects such as Stanford's Signet, Penn State's LionShare, and Internet2's own Shibboleth, setting them in the context of his unified field theory, and noted four looming issues he expects to be prominent in his talk next year: inter-federation issues, virtual organizations over P2P trust, federated and progressive (growing trust levels) PKI, and middleware diagnostics.

[Jean Pawluk](#), representing the OASIS PKI Technical Committee and PKI Action Plan coauthor Steve Hanna, presented on **Identifying and Overcoming Obstacles to PKI Deployment and Usage**. While the Technical Committee's research identified numerous obstacles to PKI deployment, the top four (Software Applications Don't Support It; Costs Too High; PKI Poorly Understood; Too Much Focus on Technology, Not Enough On Need) accounted for half the total points survey respondents assigned to indicate relative importance. The PKI Action Plan's five action items are:

- *Develop Application Guidelines for PKI Use*. This is of particular importance for the three most popular applications: document signing, secure email, and ecommerce, in that order.
- *Increase Testing to Improve Interoperability*. Again, the focus needs to be on the top three applications. Pawluk noted that smartcard implementations in particular are very vendor-dependent. She also noted the need to coordinate work so we don't have proliferating standards, which is a huge problem — bad experiences with this give people "a jaundiced view" of standards in general.
- *Ask Application Vendors What They Need*.
- *Gather and Supplement Educational Materials on PKI*. Pawluk stressed the near-complete absence of user understanding — most users have no understanding of PKI beyond "secret codes."
- *Explore Ways to Lower Costs*. Disseminating best practices is of particular interest here.

As in other sessions, prominent themes of the discussion were that technology is a much smaller part of the problem than understanding the business needs of PKI implementers and selecting tools accordingly, and that when this is done, PKI can thrive. Bill Burr observed that the math in PKI is so cool that we try to bring everything up to its standard; instead we need to figure out how people can use PKI without understanding any of the esoteric details. Rich Guida noted that he sometimes feels like he and all the people who talk about the death of PKI dwell on "different planets;" in the pharmaceutical sector in particular, the use of PKI is "blossoming." Pawluk encouraged the group to get involved in the work of implementing the PKI Action Plan, and noted that the OASIS PKI Technical Committee that's driving it (http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=PKI) usually meets via telephone.

This session was followed by a panel discussion focused on the theme: **The PKI Action Plan: Will it make a difference?** The consensus appears to be "yes, if..", with the ifs being a little different for each presenter. [Sean Smith](#) put forward his own top-three list of PKI commandments: 3) follow real-world trust flows, 2) pay proper attention to human factors, and 1) keep the goals of using PKI in mind. John Linn observed that a key question is whether deployment barriers are in PKI itself or in understanding what it can do. Most documentation is little-used and needs to

be radically simplified. Linn also stressed the importance of building in reusability across applications. Lieutenant Commander Thomas Winnenberg, chief engineer for the DoD PKI, observed that the DoD PKI has been groundbreaking in that there was no ROI concern, allowing the project to be driven by an understanding of the need for vendor-neutrality and versatility in security functions. Their big problems have been around certificate issuance, but the DoD PKI now has about four million certificates in circulation. Winnenberg stressed that the focus has to be on infrastructure — relying parties are looking to PKI for a wide variety of needs, so implementations must abstract from applications. This makes "Ask Application Vendors What They Need" a key element of the PKI Action Plan. Tim Polk stressed the importance of an iterative process of application and revision of the Action Plan. Coordination will be key (in particular liaison work with groups that don't join OASIS), as will expansion of the action items into specific, concrete tasks.

Panelist Steve Hanna asked the group for further thoughts on coordination mechanisms. Tim Polk suggested making sure that IETF meeting minutes make it to the right groups; and more generally, pushing minutes of the various groups involved out to each other, rather than relying on everyone to check up on everyone else's work. Hanna suggested that liaisons also be set up between OASIS and similar groups elsewhere in the world. Hanna also asked for thoughts on how to achieve the universally-cited goal of keeping deployment efforts focused on needs rather than technology, therefore simpler ("brass-plated," as Polk put it) whenever possible. Focusing on business needs and ROI, reusability of infrastructure across applications, and applications that make it hard for the user to do the wrong thing, were all suggested here. Russ Weiser noted that often applications are something like "I have to sign something once a year;" he suggested implementing things like this in parallel with things where security need not be as stringent but that have to be done often, like submitting a timesheet. The idea is to pick the low-hanging fruit to further usability, without worrying too much about security. With respect to reusability, Polk noted that he's become a fan of the badge/token/certificate combo — if users can't get into the building without it, they'll have it with them, and then they can use it for other things. Polk also noted that NIST has been working on a PKIX test suite and client requirements for path validation; watch <http://csrc.nist.gov>.

Conclusions

Clearly, PKI is not dead. Throughout the workshop, presenters noted the contrast between the prevailing gloomy mood and the success of their own projects. The two overarching conclusions appear to be:

1) *Understanding and educating users is centrally important.* In particular, it is crucial a) to identify the smallest possible set of things that users need to know — the things that are inherent in the nature of PKI, b) to build systems that don't require users to know more than those things, and c) to find effective ways to teach them those things.

2) *The specifics of any particular PKI deployment should be driven by real needs, and should be only as heavyweight as necessary.* The Juricert deployment is exemplary here: it was driven by the need to stop paper forgeries, avoid re-keying, and improve transaction time, and was informed by a philosophy of "PKE" — Public Key Enough.

It was in the light of this consensus that the group met to consider the future of the PKI R&D Workshop itself.

Whither PKI0x?

There was broad agreement on keeping the current R&D focus of the workshop, with particular emphases following from the conclusions above: more on human factors, and more on using the many tools available to support a wide variety of needs and architectures. With respect to the latter, attendees would like to have more of a vendor presence at the meeting — application vendors in particular. The idea would be for the PKI0x community to get a better idea of what it can do to help vendors implement the perspective developed in the course of the workshops; ideally this would become a year-round dialogue. The group would also like to hear more about international experiences and concerns,

e.g. a European report on deploying a national ID card. Finally, there was agreement that the execution of the workshop needs to be tightened up: getting proceedings out more quickly and making them more visible, and publicizing the workshop more widely and further in advance.

Non-Intrusive Identity Management

Dr. Stefan Brands

McGill School of Computer Science & Credentica

brands@cs.mcgill.ca, brands@credentica.com

March 23, 2004

ABSTRACT: This paper presents a novel architecture for digital identity management. The proposed architecture is highly secure and scales seamlessly across organizational boundaries, while at the same time protecting the privacy interests of individuals and organizations. To achieve these properties, the architecture heavily relies on Digital Credentials, a cryptographic authentication technology specifically designed to allow data subjects and organizations to securely co-manage identity-related information. We also examine the use of the new architecture in the context of three emerging information-sharing applications: Electronic Health Record management, E-Government, and Digital Rights Management.

1. Introduction

Most people are registered in many hundreds if not thousands of databases scattered across disparate systems. In identity management jargon, individuals have multiple *network identities*: collections of information that relate to an individual, that are created and managed as single units in a network, and that are stored in electronic form. Advancements in networking technologies make it increasingly easy to collect and collate these network identities.

Of course, this cross-domain aggregation power by itself is not of much value to organizations, unless it is combined with the ability to determine which network identities correspond to the same individual. Traditionally, identifiers such as health insurance numbers and Social Security Numbers serve as keys to facilitate such cross-linking. The current efforts in the electronic world to enable cross-domain identity management and information sharing rely on their own unique cross-domain identifiers, such as biometric templates and digital certificates.

For businesses, an increase in cross-domain linking power ultimately translates into increased sales and cost reduction. For government organizations, the ability to share client information translates into more efficient interactions with citizens and an improved ability to detect and contain fraud. Individuals stand to benefit as well from these, at least in principle.

Privacy and security concerns

The increased introduction of (and reliance on) cross-domain identifiers also brings serious privacy risks. Target marketing can turn into spamming, service customization can turn into unfair price discrimination, hackers and insiders can cause systemic denial of access to targeted individuals, and so on. For these and other reasons, many people provide false identity information when accessing on-line services.

Indeed, the business goal of cross-domain digital identity management is directly at odds with the privacy interests of individuals. What businesses since a few years refer to as network identity is essentially what data protection legislation around the world already since the early eighties refers to as *personal information*: information about a “data subject” whose identity can reasonably be ascertained from the information. Data protection legislation requires organizations to protect personal information in accordance with several privacy principles, one of which is information security safeguards.

Intra-enterprise security needs, however, are much lower than cross-domain data protection requirements. Indeed, while password-only authentication is often adequate for internal access to organizational resources, in cross-organizational contexts it would give outside organizations unacceptable impersonation powers. In the context of cross-domain access management, traditional information security products (such as firewalls, anti-virus software, intrusion detection systems,

and vulnerability assessment tools) break down as well; with trust domains being logical rather than physical, security must be tied to the data itself rather than to the perimeter of its repository.

In short, organizations are starting to discover that the arsenal of security tools they use for intra-organizational data protection is not appropriate to protect information that is shared across organizational boundaries.

Federated identity management

The currently prevailing industry approach to address this situation is to centralize all the authentication power from different domains into a single trusted domain that acts on behalf of its constituent organizations. With *federated identity management* architectures, such as those pushed forward by Liberty Alliance, organizations do not authenticate access requestors themselves, but instead query a trusted Identity Provider that does the authentication for them. The Identity Provider simply returns an authentication assertion as to the validity of the identity claim of the access requestor, which the relying organization uses in its own authorization process. This approach in effect maps the cross-domain context back to the traditional single-domain context, which organizations know how to handle using traditional authentication techniques, be they password-only authentication, Kerberos, or perhaps PKI. (Indeed, PKI vendors generally consider federated identity management, and notably standardization efforts such as SAML, as what will rescue PKI from an untimely death, since a full-fledged certificate infrastructure is unnecessary.)

However, centralizing systems of an inherently decentralized nature brings its own administration, scalability, security, and privacy problems, which may be far worse than the original problem one was seeking to solve. In its original Passport architecture, for example, Microsoft relied on the centralization of all authorization data, and was forced to back down following complaints from consumer groups, EU officials, and organizations that were reluctant to entrust Microsoft with their customer data. The Liberty Alliance proposal improves over the original Passport scheme by leaving personal data at the organizations that collected it, but the authentication power (and therefore the ultimate access control power) remains centralized within each circle of trust.

At its core, federated identity management architectures such as the Liberty Alliance proposal and Microsoft's revised Passport scheme are centralized authentication architectures. Indeed, the Identity Provider's role in the Liberty Alliance architecture greatly resembles that of Visa or Mastercard among their respective "circles" of merchants: within its circle of trust, the Identity Provider can track, trace and link in real time all the interactions between users and organizations. (It may not know the transaction details itself, but that by no means is enough for information privacy.) The Identity Provider can even impersonate users and falsely deny them access everywhere. Furthermore, Identity Providers are highly appealing targets for fraudulent insiders and hackers. On top of that, relying organizations do not get the strength of the authentication mechanism used by the Identity Provider, but merely that of the session maintenance mechanism used when redirecting the user between the organization and the Identity Provider; impersonating a user or cloning access privileges depends merely on the difficulty of getting to a session cookie, rather than on the difficulty of getting to the user's secret key (which could be stored on a smartcard). For an in-depth analysis of the Liberty Alliance architecture, see [1].

More generally, the privacy, security, and scalability problems of centralized authentication architectures have been well-documented in the past two decades by the professional cryptography and security community. In the context of "unbalanced" B2B digital identity management (where organizations inherently place asymmetric trust in a central party), the shortcomings of industry's current federated identity management efforts may not be problematic. Beyond that, however, they may well turn out to be a showstopper. Collaborative enterprise efforts, where participating organizations are equals ("balanced" B2B), may already prove too much of a stretch, not to mention G2C, B2C, and C2C applications.

The need for new approaches

The growing mismatch between the security needs of cross-domain identity management and traditional security tools and practices is not all that surprising. The currently prevailing authentication techniques (password-only, biometrics, Kerberos, PKI) were all invented more than two decades ago, when open networks were hardly existent, let alone organizations seeking to securely share identity-related information over such networks. At that time, privacy legislation was virtually non-existent. The only privacy protection that the designers of the traditional security techniques had in mind was protection against unauthorized outsiders (e.g., wire-tapping). In the new frontier of cross-domain access and identity management, however, the biggest threats to privacy do not come from outsiders, but from insiders.

To better understand the shortcomings of PKI and other authentication mechanisms that were not designed with cross-domain identity management requirements in mind, it is important to understand the relation between (information) security and privacy. *Security* is generally defined as the extent to which information can be stored and transmitted in such a manner that data access is limited to authorized parties. *Privacy* is generally defined as “the claim of individuals to determine for themselves when, how, and to what extent information about them is communicated to others.” In accordance with the Fair Information Principles of the OECD (which form the basis of most of today’s data protection legislation around the world), “security safeguards” is only one of the eight principles necessary to achieve privacy. In contrast to security, which is aimed at preventing access by unauthorized outsiders, the other basic privacy principles are primarily aimed at unauthorized use by insiders. As such, security safeguards are necessary to achieve information privacy, but not sufficient. Ironically, traditional authentication technologies have a highly adverse impact on two of the most important privacy principles: collection and use limitation. They are, in fact, privacy-invasive technologies.

What this paper is about

Two decades of research in modern cryptography has shown that security and privacy are not trade-offs, but that they are mutually reinforcing when implemented properly. A fundamental premise of modern cryptography is that the need to rely for privacy on Trusted Third Parties (such as the Identity Provider in federated identity management) can be eliminated. This brings us to the goal of this paper: to present a non-intrusive identity management architecture that is highly secure and that scales seamlessly across organizational boundaries.

2. Non-Intrusive Identity Management

Before describing the proposed architecture, we give an overview of the state-of-the-art authentication primitive that is at the core of the new approach to cross-domain identity management: *Digital Credentials*. Our architecture will rely on this new primitive in four ways.

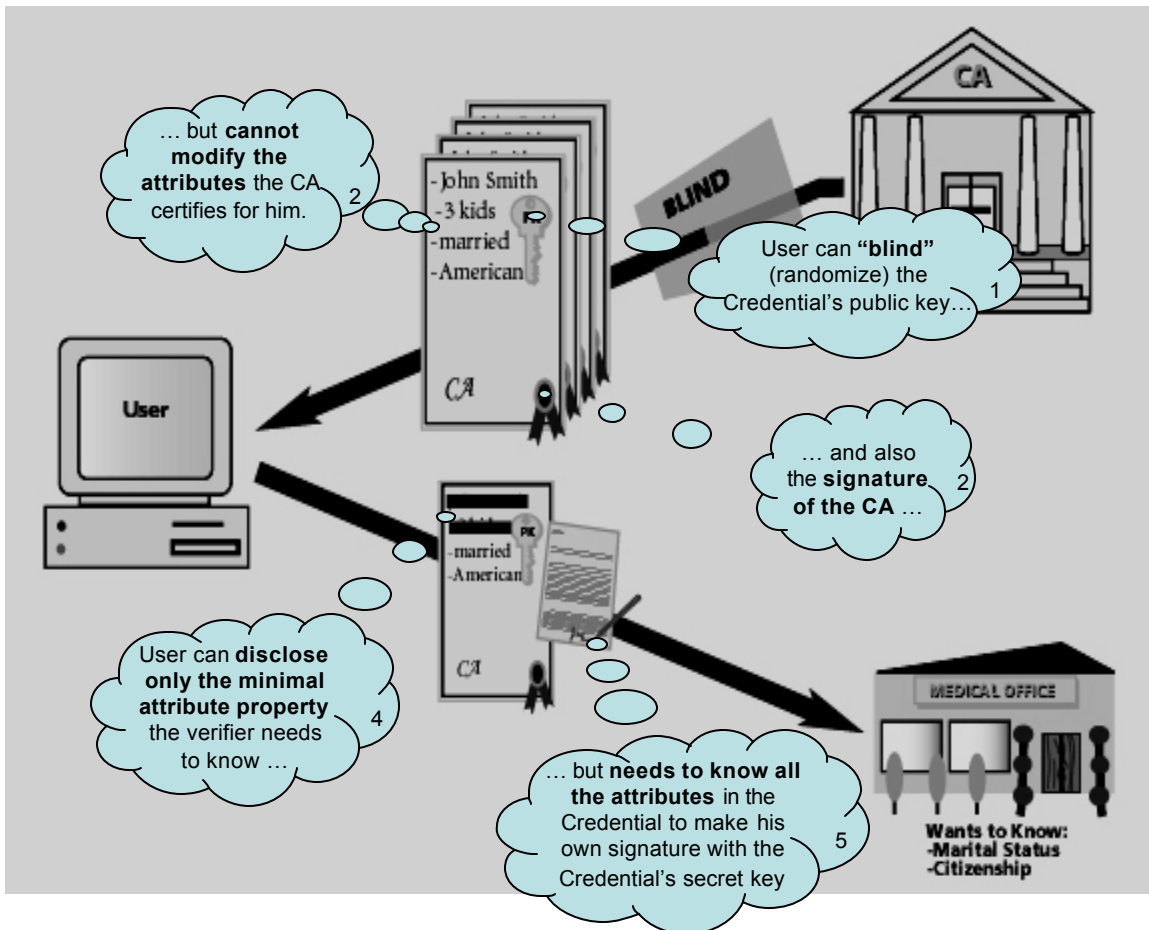
Digital Credentials

Digital Credentials are the culmination of two decades of scientific advances by dozens of professional cryptographers, starting in the early eighties. They are basic cryptographic constructs, much like digital signatures and equally efficient, but with much greater functionality. Specifically, Digital Credentials provide *fine-grained privacy control* at every step in the life-cycle of certified personal data that is being sent around. As well, they have security properties that go well beyond what can be achieved for X.509 identity and attribute certificates (in ways that may at times seem counter-intuitive), and can be implemented in low-cost smartcards without cryptographic coprocessors. They support all the traditional authentication strengths, from software-only protection to military-grade two-factor and three-factor security. Technically, Digital Credentials are issued and shown as follows:

- **(Issuing protocol)** Digital Credentials are issued to applicants by Credential Authorities. Each Credential Authority has its own key pair for signing messages. When issuing a Digital

Credential to Alice, the issuing Credential Authority through its own digital signature binds one or more attributes to a Digital Credential public key, the secret key of which only Alice knows. (An attribute can be any information.) The entire package that Alice receives is called a Digital Credential. Although the sequences of zeros and ones that make up Alice's public key and the signature of the Credential Authority are unique for each Digital Credential, the Credential Authority cannot learn who obtains which sequences; they are *blinded* during the issuing process. At the same time, the blinding operations that Alice can perform are *restricted* in such a manner that Alice cannot modify the attributes that the Credential Authority encodes into her Digital Credential. What's more, some or all of the attributes in Alice's Digital Credential could initially be provided to the Credential Authority by Alice herself, by her smartcard, or by another organization, without the Credential Authority being able to learn them.

- **(Showing protocol)** To show her Digital Credential to Bob, Alice sends her Digital Credential public key and the signature of the Credential Authority. She also digitally signs a nonce, using her secret key. (A nonce is a random number, the concatenation of Bob's name and a counter, or any other fresh data provided by Bob.) Bob cannot replay Alice's information for his own benefit in another transaction, since in each showing protocol execution a new nonce must be signed; this requires knowledge of Alice's secret key, which never leaves Alice's device. At the same time, Alice can *selectively disclose* to Bob a Boolean property of the attributes in her Digital Credential (this goes well beyond what can be done with a paper-based certificate and a marker), while hiding any other information about them. Importantly, however, it is infeasible for Alice to demonstrate any property without her actually knowing all the attributes encoded into her Digital Credential (including those that she does not disclose). To convince Bob that the claimed property is true, Alice's signature on Bob's nonce doubles up as a proof of correctness.



A detailed description of how these basic properties are achieved in a highly practical manner is outside the scope of this paper. A technical overview of Digital Credentials can be found in [2], and the full details appear in [3]. As explained in these references, by carefully exploiting these basic properties of Digital Credentials, one can efficiently realize all of the following features:

- **(Privacy of Credential holders)** Digital Credentials accommodate fully adaptable levels of privacy ranging from user-driven anonymity to government/enterprise-mandated identification. They support automated negotiation of credential information, ensuring the disclosure of only the minimum credential information needed to meet the authorization requirements of an access provider; this minimizes the risk of identity theft, and preserves privacy. The selective disclosure technique can be applied not only to attributes encoded into a single Digital Credential, but also to attributes in different Digital Credentials, possibly certified by different Credential Authorities.¹ There is no need to trust third parties to protect one's privacy: even if all the parties that rely on Digital Credentials actively conspire with all Digital Credential issuers and have unlimited computing resources, they cannot learn more than what can be inferred from the assertions that Digital Credential holders willingly and knowingly disclose.
- **(Privacy of Credential verifiers)** In many situations, verifiers may want or need to pass on Digital Credential evidence to central parties (e.g., for online revocation status checking, to enable fraud detection on behalf of multiple access providers, to allow statistical data gathering, or to serve as transaction receipts). A Digital Credential verifier can *selectively hide* any or all of the information that a Digital Credential holder selectively disclosed to it, before forwarding that Digital Credential. In other words, the verifier can forward non-repudiable transaction evidence that proves to third parties no more than exactly what it wants the evidence to prove; this may be much less than what the Digital Credential holder selectively disclosed to the verifier. By way of example, consider a patient-physician interaction or a consumer-merchant transaction; while the customer may have no problem identifying himself to his doctor or to the merchant, the latter parties may not want to disclose their customer's identity to third parties.
- **(Strong accountability)** Digital Credentials offer audit capability for non-repudiation and to assess compliance with regulatory requirements, through digital audit trails and receipts that facilitate automated dispute resolution. Malicious parties, including Credential Authorities, cannot frame a Digital Credential holder by making it look as if he or she participated in a transaction, even if they would collude and would have unlimited computing power or special knowledge of trapdoor information. Audit trails can be kept in the form of role-based digital signatures; in case of a dispute, the transaction originator cannot disavow the origin.
- **(Pooling protection)** Different people can be prevented from pooling together multiple Digital Credentials in order to enjoy access privileges that they would not enjoy on their own. Hereto the access provider requires the access requestor to demonstrate that any Digital Credentials that he or she shows contain the same built-in identifier. Owing to the selective disclosure property, an honest Digital Credential holder can demonstrate this without disclosing the built-in identifier.
- **(Discarding protection)** Digital Credentials can be used to prevent the discarding of authenticated information that an access requestor would rather not show. A mark for drunk driving, for instance, can be tied into a driver's license Digital Credential that specifies that the holder is authorized to drive. Once again owing to the selective disclosure property, the owner can hide the mark whenever it need not be disclosed.
- **(Lending protection)** Lending of credential information can be discouraged by wrapping the information into a Digital Credential and encoding confidential data of the legitimate owner into it. The legitimate owner can hide this data (again owing to the selective disclosure property), but

¹ Rather than encoding many attributes into a single Digital Credential, it may be preferable to distribute them across multiple Digital Credentials. This helps avoid the aggregation of an individual's attributes by a single Credential Authority, improves efficiency when many attributes need to be encoded independently, and removes the need to update certificates more frequently than otherwise needed.

the Digital Credential cannot be used without actually knowing the confidential data. (Note that this measure does not rely on credential holders using tamper-resistant devices.)

- **(Dossier-resistance)** A Digital Credential can be presented to an organization in such a manner that the organization is left with no evidence at all of the transaction (much like showing a passport without letting the other party make a photocopy) or such that the verifier is left with self-authenticating evidence of only a part of the disclosed property. Furthermore, the self-authenticating evidence can be limited to designated parties. In case of a dispute, the disclosed property can always be revealed in full.
- **(Limited-show credentials)** A limited-use Digital Credential can contain a built-in identifier, value token, or self-signed fraud confession, that will be exposed if (and only if) the Digital Credential is shown more than a pre-authorized number of times.² These *limited-show* Digital Credentials (which can be used to design the digital equivalent of stamps, coins, tickets, and so on) have no obvious paper-based analogue. The limited-show property holds even when Digital Credential holders are free at each occasion to choose the attribute properties that they demonstrate, and even if they conspire with verifiers (who, as mentioned, are able to hide any information disclosed to them before forwarding transaction evidence). Limited-show Digital Credentials are highly practical: to be able to compute a built-in identifier in case of fraud, a footprint of a mere 60 bytes must be stored for each Digital Credential shown, regardless of the complexity of the property disclosed and regardless of the number of encoded attributes.
- **(Negative authentication)** This property allows the holder of a Digital Credential to demonstrate that he or she is *not* someone listed on a blacklist, without enabling identification. More generally, the holder of a Digital Credential can demonstrate that the data in the Digital Credential does *not* meet certain conditions, without revealing more.
- **(Recertification and updating)** In many cases the right to access a service comes from a pre-existing relationship in which identity has already been established. An individual can present a certified public key for recertification or for updating to a Credential Authority, without enabling it to learn the current values of the attributes in the Digital Credential. Of course, the Credential Authority could require the individual to demonstrate an attribute property before certifying the Digital Credential or its updated version.
- **(Information can reside anywhere)** Digital Credentials can be held both locally (on a device of the user) or remotely, and can be managed using roaming. In the extreme, organizations can do away entirely with central databases containing sensitive personal information, by securely distributing each database entry to the individual to whom it pertains; the unique security properties of Digital Credentials ensure that unauthorized users cannot modify, discard, pool, or lend their own credential information, nor can they prevent it from being updated (without locking themselves out of the entire system).
- **(Smartcard Implementation)** Digital Credentials can be issued to, or embedded in, smartcards and other tamper-resistant devices; this provides a second layer of protection (on top of the abovementioned cryptographic protections) against loss, theft, lending, pooling, copying, and discarding of Digital Credentials. As well, the Digital Credential holder's smartcard can prevent other kinds of unauthorized behavior by its owner, and can protect him against "virtual" extortion attempts. The storage and computational burden for the tamper-resistant device can be off-loaded almost entirely to another user device that need not be tamper-resistant (such as a handheld device, a laptop, or another chip on the same smartcard that need not be trusted by the system provider), while preserving all of the smartcard's security benefits; literally billions of Digital Credentials can be securely managed in this manner using a single 8-bit smartcard chip without a cryptographic co-processor.

² Alternatively, copying and reuse can be prevented by resorting to online Digital Credential validation by a central party, but this may pose a serious performance bottleneck.

- **(Secure multi-application smartcards)** Smartcards can be used as multi-application devices, without introducing any of the privacy and security problems caused by other technologies. Specifically, different application providers can all share the same secret key stored in a user's smartcard to derive the security benefits of that smartcard. The certificates will have uncorrelated secret keys which cannot be determined by anyone including the smartcard supplier, and all Digital Credentials can be revoked separately. The application software on the user's trusted computer ensures that smartcards attacks and data leakages are impossible. Moreover, different applications relying on the same smartcard can be fire-walled through the application software running on the patient's trusted computer, rather than the application providers and the card holder having to trust the smartcard issuer.
- **(Managed security services)** With an increasing number of incompatible authentication mechanisms in use, organizations that need to make authorization decisions will increasingly ask trusted authorities to issue and/or verify the credential information presented by their clients. With Digital Credentials, Credential Authorities can certify sensitive information on behalf of organizations without being able to learn that data, and Revocation Authorities can validate certificates (using OCSP or other standards) without being able to learn the identities of the clients of organizations (even when these clients disclose their embedded identities to the organizations they transact with). In this manner, organizations can outsource core tasks related to digital authentication and authorization, without having to provide their managed security services provider with competitive data or customer information for which they could incur legal liabilities. Even the role of the tamper-resistant smartcard can be outsourced, removing the logistical problem of securely distributing tamper-resistant devices.³

With this set of features in mind, we are now sufficiently prepared to discuss our approach to cross-domain digital identity management.

Identity management based on Digital Credentials

We will refer to our proposed architecture as the Credential Management Platform (CMP). CMP is characterized by three central notions: records, participants, and protocols.

A *record* is a logical collection of information. Records may be held in a central database, may be distributed across multiple databases, or may be held locally on a user device. In the first two cases the record is called a Remote record; in the latter case it is called a Local record. In general, Local records offer greater security and privacy to access requestors, but may be less convenient. Implementations of CMP could facilitate the automated sharing and synchronization of Local and Remote records in accordance with application-specific administrative data, to allow multiple records to be managed electronically as one logical entity.

A record contains two kinds of information:

- **Attributes:** An attribute is any personal data, corporate intelligence data, or otherwise sensitive information to which access must be guarded. Attributes may be encrypted by a key known only to a participant; this is useful for instance when attributes that are normally held in a Local record are temporarily stored on a public network to support roaming access by other devices.
- **Related administrative data:** The administrative data describes rules that specify by whom each attribute in the record may be read, written, modified, or otherwise accessed. Administrative data can include audit trails (possibly digitally signed) for access events.

CMP distinguishes between three kinds of attributes in a record:

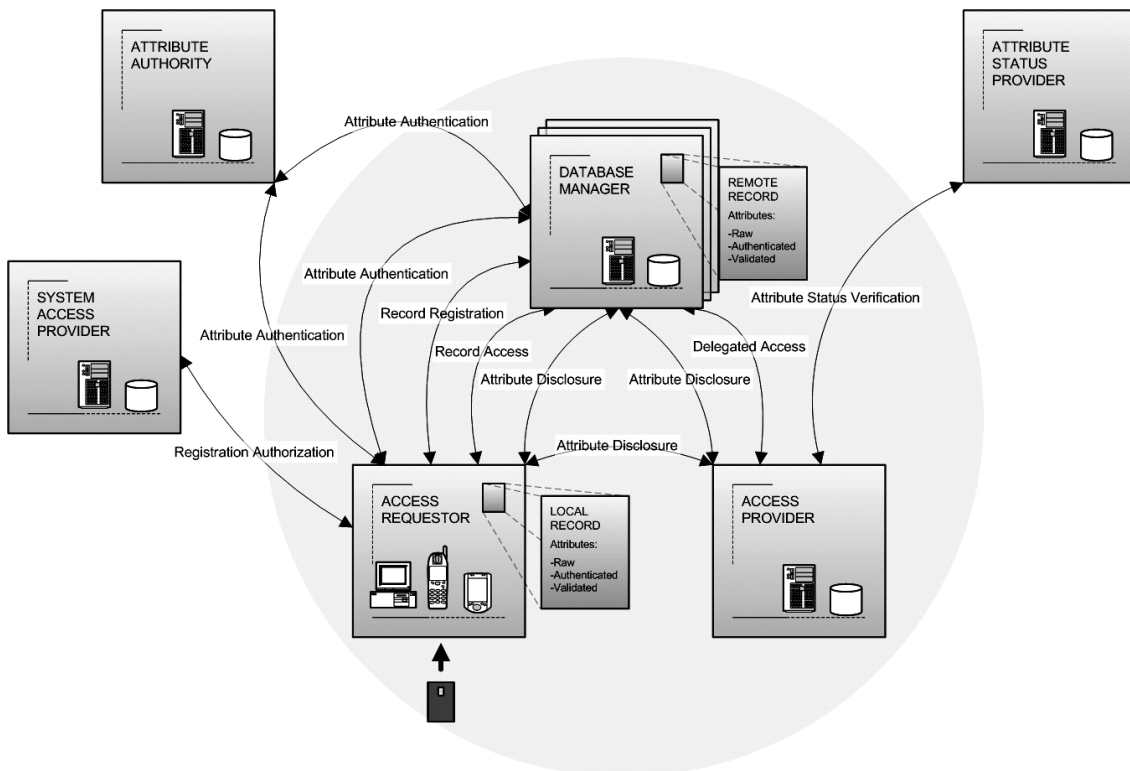
- **Raw attributes:** These are attributes specified by any party without any guarantee as to their validity. Personalized display or content preferences for a Web site are an example. Modification

³ Although every transaction of a Digital Credential holder will now require the real-time involvement of a third party that guarantees protection of the user's secret key, that third party cannot learn any details that could lead to a privacy compromise (other than knowing the transaction times of pseudonymous users).

or discarding of raw attributes by unauthorized participants might cause inconvenience to the party to whom the data pertains, but would not adversely affect the security of any other party.

- **Authenticated attributes:** Attributes that are digitally authenticated by a participant by means of a digital signature, but without prior verification of their validity. This prevents other participants from modifying the attribute. In an on-line chat group discussing gender-related issues, for example, a person might wrongly specify his own gender but would be stuck with it in future sessions.
- **Validated attributes:** Attributes that are digitally authenticated by an “Attribute Authority” only after the validity of the attribute has been verified by that Attribute Authority.⁴

Authentication of authenticated and validated attributes takes place by wrapping one or more attributes into a Digital Credential, to offer unique security, privacy, and usability benefits. Different attributes may be packaged either into separate Digital Credentials or into the same Digital Credential. Furthermore, different Attribute Authorities may authenticate the same attribute by packaging it in different manners. By way of example, consider an electronic patient record: multiple doctors may digitally “sign off” on the same entries in a patient record. More generally, multiple Attribute Authorities may package the same or different overlapping subsets of attributes in a record in different ways into Digital Credentials. In this manner, access providers can be assured that the data entries on which they rely have been entered by authorized parties, and different parties can effectively maintain partial ownership of information in a record. Not even the party (or combination of parties) controlling the storage of a record can modify, delete, or add information, unless they are properly authorized.



A *participant* is a device or application (or a collection of devices or applications) that acts either autonomously or on behalf of an individual, a group, or an organization. For simplicity we will interchangeably refer to participants as both devices or applications and the parties they represent. CMP distinguishes between six types of participant:

⁴ Attribute information may be supplied to Attribute Authorities by “Registration Authorities” who are responsible for validation; we do not explicitly show Registration Authorities in our architecture, however.

- **Database Manager:** A party that controls the physical storage of records.
- **Attribute Authority:** A party that issues authenticated or validated attributes. These attributes may be valid only a limited number of times or only for a limited-time period.
- **System Access Provider:** A special Attribute Authority responsible for granting participants the right to “initialize” Remote records (and possibly to subsequently manage it in a co-owner role). The System Access Provider issues Registration tokens, either one per participant until expiry of the token or a new one at regular time intervals or when requested.
- **Access Requestor:** A party interested in accessing a service that requires an authorization decision. The Access Requestor may be represented by a PC, a handheld device, a mobile phone, a smartcard, or any other device capable of computing and communicating.
- **Access Provider:** A party that relies on some or all of the attribute information in a record in order to make an authorization decision pertaining to an Access Requestor. Attributes in the record (more generally, properties about attributes in one or more records) are presented to the Access Provider either by the Access Requestor or by the Database Manager. In the latter case, either the Access Requestor's active involvement or prior explicit consent (in the form of a Delegation token) is needed. The Access Provider may resort to an Attribute Status Provider to complete the verification of authenticated and validated attributes.
- **Attribute Status Provider:** A party that verifies the status of one or more attribute-related requests presented by an Access Provider. Its primary role is to verify the revocation status of validated Attributes, to manage and issue updates of revocation lists, and to keep track of the number of times a limited-show attribute has been used.

In a real-world application, there will normally be many instantiations of most types of participant. For instance, in an electronic health record management system, each doctor authorized to update patient records would be an Attribute Authority. Of course, the roles of multiple participants from the same or from different systems may in practice all be performed by the same party.

Participants interact with each other by means of *protocols*. CMP distinguishes between seven basic protocols:

- **Registration Authorization:** A protocol between a System Access Provider and an Access Requestor whereby the Access Requestor obtains a Registration token allowing him to subsequently initiate a Remote record. The Registration token may be issued to a tamper-resistant device (e.g., a smartcard) of the Access Requestor for greater security.
- **Record Registration:** A protocol between an Access Requestor and a Database Manager whereby the Access Requestor presents a Registration token to initialize a record. As part of the protocol, the Access Requestor and the Database Manager specify administrative data.
- **Record Access:** A protocol between an Access Requestor and a Database Manager whereby the Access Requestor accesses a record stored by the Database Manager in order to read, write, or modify attribute information. The Access Requestor must show an Authorization credential (which may be the Registration token) to demonstrate proper access rights.
- **Attribute Authentication:** A protocol whereby an Attribute Authority issues an authenticated or validated attribute for entry into a Local or Remote record. The Attribute Authority issues the authenticated attribute either upon receiving an Authorization credential or upon receiving authenticated attribute information issued by another Attribute Authority.
- **Attribute Disclosure:** A protocol whereby an Access Requestor discloses attribute information to an Access Provider. The protocol can be conducted either with or without the assistance of the Database Manager. For Local records, there is no need to involve a third party in order to disclose attribute information to the Access Provider. For Remote Records the Access Requestor can disclose the attribute information to the Access Provider either by directly retrieving it online

and forwarding it to the Access Provider, or by routing its own access request through the Access Provider to the Database Manager.

- **Delegated Access:** CMP allows the Access Requestor to provide the Access Provider with a digitally authenticated Delegation token specifying the latter's access rights, so that the latter can later on access a record (perhaps for a limited period of time or a limited number of times) without further involvement from the Access Requestor's side.
- **Attribute Status Verification:** A protocol between an Access Provider and an Attribute Status Provider whereby the Access Provider requests and obtains information on the status of an attribute beyond what it can infer from the attribute itself. Attribute Status Verification may take place either on-line (in conjunction with an Attribute Disclosure protocol) or off-line. For short-lived authenticated and validated attributes, the Attribute Status Verification protocol may not be needed.

All tokens, access requests, and other forms of authentication in CMP are implemented using Digital Credentials. Specifically, CMP relies on Digital Credentials in four basic manners:

- To implement access privileges, entitlements, delegations, and any other attributes that access requestors show to access providers to allow them to make local authorization decisions;
- To implement privacy-enhanced digital identity certificates (usable as digital pseudonyms where identification is not required) that allow the separation of different spheres of activity;
- To authenticate the entries of electronic records stored in central or distributed databases; and
- To implement digital audit trails and digital receipts that witness details of access requests.

Unique benefits of CMP

As a direct consequence of using Digital Credentials throughout the CMP architecture, a number of unique benefits arise, including the following:

- The Registration token can be presented in a manner that does not enable identification of the Access Requestor. (Digital Credentials encompass identity certificates as a special case: an identifier is just one of infinitely many attributes that can be encoded into a Digital Credential, and the Digital Credential holder can disclose it whenever desired.)
- For Remote records, the Access Requestor can choose to be identified or to remain pseudonymous. The ability to pseudonymously hold a Remote record reduces the risk of identity fraud, and minimizes the damage that can be done by malicious insiders and outside attackers. In the case of a dispute a pseudonymous Access Requestor will not be able to deny having accessed the record; only pseudonymous Access Requestors who did not access the record can prove they did not do so.
- In the case of Local records, CMP allows the Access Requestor to be fully anonymous. The authenticators of attributes in the record can strongly discourage the Access Requestor from cloning or lending his attributes. Furthermore, the Access Requestor can present the Attribute Authority with a previously issued authenticated or validated attribute in order to have it re-authenticated or updated, without enabling the Attribute Authority to learn more than it strictly needs to. In the case of a limited-show attribute, a built-in identifier, value token, or self-signed fraud confession will be exposed if the attribute is used more times than allowed.
- The Access Requestor can disclose only the minimum attribute information (such as a particular property of multiple attributes) needed to meet the authorization requirements of the Access Provider. (In case the attribute is stored in a Remote record, this requires the Access Requestor to have some trust in the Database Manager.)
- Access Providers that know an Access Requestor under different unlinkable pseudonyms can enable the Access Requestor to transfer authentication attribute information from one pseudonym to another without creating pseudonym linkage, while at the same time preventing

the Access Requestor from showing attributes that belong to another Access Requestor (even if Access Requestors collude).

- In case Digital Credentials are issued to smartcards, all computationally expensive operations for the smartcard can be off-loaded to a more powerful device; virtually no smartcard storage space is required in that case, so that plenty of room is left for a software solution to protect against sophisticated attacks such as differential power analysis. Also, CMP can offer protection against fake-terminal attacks and smartcard data leakage by routing communications from and to the smartcard through a device trusted by the card holder.
- Attribute Authorities can digitally authenticate information on behalf of others without being able to learn attribute data that they have no need to know. Likewise, Attribute Status Providers can validate certificates without being able to learn the identities of access requestors and access providers. In this manner, Access Providers and Database Managers can outsource core tasks related to digital authentication and authorization to security specialists, without having to provide them with sensitive information.
- Attribute information can be presented to the Access Provider in such a manner that the Access Provider is left with self-authenticating evidence that proves only a part of the Attribute property disclosed by the Access Requestor; this enables the Access Provider to pass on the evidence to third parties (such as the Attribute Status Provider), while protecting its own privacy, complying with privacy legislation, and avoiding leakage of competitive intelligence.
- For Access Providers, Record Access can be identified, pseudonymous, or anonymous. The latter two cases prevent the Database Manager or the Attribute Status Provider from gaining competitive intelligence on Access Providers or from improperly rejecting valid requests for access on the basis of the identity of the Access Provider. At the same time, the Access Provider can disclose exactly that which is required to enable the Database Manager to make its own authorization decision: CMP provides for role-based access. The Database Manager and other parties can strongly discourage the Access Requestor from reusing, lending, pooling, discarding, or cloning his access rights, even for pseudonymous access.

3. Example Applications

We now discuss the benefits of using CMP in the context of several emerging applications that fundamentally rely on cross-domain identity management.

Electronic health record management

An Electronic Health Record (EHR) is defined as the health record of an individual that is accessible online from many separate, interoperable automated systems within an electronic network. EHRs can contain a variety of data and can be used for different purposes by different parties involved in health care. The grand vision of EHR infrastructures is the interconnection and reusability of all recorded health information, regardless of where it is stored, so that all relevant health information can electronically flow to wherever it is needed.

Nothing will become of this vision, however, unless critical privacy and security problems are overcome. Studies reveal that most patients do not trust the administrators of national health services and other insiders in the health care system with the control over their personal health information. Often, their trust does not extend beyond their own care providers, and indeed the opportunities for privacy invasions due to secondary use of health record information are enormous. Organizations with a justified need (according to current widespread regulations) to access health information include government and private health plans, insurance companies, administrators, hospitals, doctors, pharmacies, employers, schools, researchers, data clearinghouses, accreditation and standard-setting organizations, laboratories, pharmaceutical companies, practice management system vendors, and billing agents.

Privacy is also sought by medical practitioners. Many doctors do not like the idea of central parties (such as health insurance organizations) being able to monitor all their actions, since they feel this negatively impacts their autonomy; in many situations, they would prefer to be able to access information on the basis of their role rather than their identity, and they certainly do not want identifiable digital evidence of all their interactions with patients to automatically flow to central parties. Role-based access is also preferred by medical researchers, for accessing online disease registers and other medical databases.

With CMP, an EHR is simply a Local or Remote record, or the logical combination of several such records. Attribute Authorities are health care professionals and possible other entities (including the patient himself) who add digitally authenticated statements to EHRs. EHRs can be securely managed by both the data subject and his health care professionals, in a manner that simultaneously protects the data subject's privacy interests, the professional's liability interests, and the legitimate interests of researchers and other third parties:

- Each patient can co-manage his health information together with selected physicians. A record can be managed electronically as one logical entity, even though different parts may reside in different physical locations. Each party with access rights can be assured that the data entries on which it relies have been entered by authorized parties, through either role-based or identity-based digital signatures. In this manner, health care service providers can effectively maintain partial ownership of a data subject's health information.
- By providing patients with tamper-resistant smartcards, health care providers can maintain even greater control over their own contributions to EHRs, since the cards can further limit the ability of patients and others to manipulate entries. Literally billions of authenticated EHR entries (possibly originating from different health professionals) can be securely managed using a single 8-bit smartcard. Cards can be issued to patients by a central entity that cannot compromise the legitimate privacy and security interests of patients and health care providers that ride along on the added security provided by the card.
- At the same time, patients as well as health professionals are able to selectively disclose authenticated health data in anonymous or pseudonymous form (with or without certifications). Patients can also delegate the right to do so to their doctors (e.g., to over-ride protections in emergency situations) or to third parties (e.g., for research purposes).

CMP in effect creates a continuum between health records maintained by health professionals and health records maintained by data subjects, seamlessly unifying the two approaches and covering the entire spectrum of possible rights management settings. In the CMP approach, the issue of where the health data resides hardly matters anymore; it is all about who has electronic access to which parts of a record.

E-government

E-government refers to the electronic delivery of government services to citizens. The primary objective is to simplify the interaction with citizens and institutions. In the past three years, many municipal, provincial, and federal governments around the world have established an on-line presence. Among the leading countries to bring government online are the United Kingdom, Canada, and the United States. Market analysts distinguish between five phases of e-government: (1) providing information via Web sites; (2) electronic service delivery; (3) improving operations through Web interfaces and electronic data exchanges; (4) moving toward more personalized electronic service delivery ("e-CRM"); and, (5) introducing Web-based collaborative technologies. Implementing CRM initiatives is widely considered a key priority to provide personalized citizen self-service.

In most cases, government organizations will need to be able to securely make authentication and authorization decisions about citizens who request electronic access to their services. Liberty Alliance is already being viewed with increasing interest by e-government architects. Indeed, the considerations of government for managing identity-related information in part match those of

industry. Governments however have a stronger interest in protecting the security and privacy of individuals and private sector organizations, and for good reasons. For instance, an August 2000 survey by Hart-Teeter about U.S. citizens' view of e-government services found that 53% of respondents were extremely concerned with the potential loss of privacy, and in a Gartner survey in 2001, nearly 70% of consumers cited privacy concerns as one reason that could make them stop using e-government services. As well, it would be most awkward for government not to live up to the spirit of its own data protection legislation, and ultimately the stability of democracy may be put on the line if a privacy-threatening infrastructure would be implemented.

On the security side, progress is being made in the right direction. Indeed, according to the Giga Information Group in June 2002, "in some technologies, like smartcards, biometrics and electronic records management, the government is ahead of business." Many governments are keen on access management systems based on smartcards, not only for citizens but also for its own employees and to replace driver's licenses, airport security documents, passports, and so on. On the privacy side, however, governments are struggling. Consumer outcry, trade group complaints, potential violation of privacy laws, and complaints by data protection commissioners have already lead to the suspension of several national PKI e-government initiatives.

Using CMP, it is easy to see how security, scalability, privacy, and general performance requirements can be reconciled. Consider the case of personalized access to on-line government services. A user would retrieve digital pseudonyms in batch from a central certificate issuer. The user would register a different pseudonym with each on-line service provider, which the service provider would link to its own program identifier for that user (following its own one-time authentication of the user's program-specific identity, or following an "introduction" by another organization). Due to the unlinkability of pseudonyms, government service providers do not gain cross-domain profiling powers that were not present in the legacy system. The certificate issuer can serve as a managed security services provider to government organizations, by providing some or all of the security features described previously. At the same time, the certificate issuer can be prevented from gaining any tracking and tracing powers, and can even be prevented from learning the identities of the certificate requestors as they retrieve pseudonyms. By embedding a unique "identifier" (e.g., a random number) into all of a user's pseudonyms, the certificate issuer can ensure that users can transfer certified personal information from one government organization to another without users being able to lend or pool personal information; in this manner, government organizations can reliably share user information without obtaining cross-domain profiling powers. As well, the certification of identity information by each government organization can be delegated to the certificate issuer without the latter being able to learn the information itself.

Digital rights management

Digital Rights Management (DRM) is generally defined as the collection of tools and technologies for protecting copyrights and other rights on digital media. DRM is an umbrella term: no single tool or technology suffices to guarantee access and content usage controls throughout a digital content distribution infrastructure.

DRM deals with authorization decisions about access to resources, and as such it is an application of access management. However, DRM places stronger requirements on fraud prevention than general access management. Namely, access management in general does not deal with long-lived access, while DRM also seeks to control usage by authorized users after they gain access to a resource.

All modern DRM systems have at their core the notion of a digital license, and most deal with content and licenses in a separate manner, along the following lines:

- Licenses are issued when access is requested, while content is made freely available in encrypted form to prevent access by unauthorized parties. To access protected content, the client must obtain a digital license that specifies how the content may be used.

- To consume protected content, the client connects to a clearing house and requests a digital license for the content. The request requires the client to send a unique identifier that identifies him and/or a specific client device that will play the content. The request is typically initiated by the client's software application or hardware device upon the client's first attempted access.
- Assuming the clearing house makes a favorable authorization decision for the client, it sends the requested digital license to the client. The client's device or application, which is presumed to be secure against tampering, then decrypts the license and displays or otherwise makes available the content to its user in accordance with the usage rules.

By separating content from licenses, content providers can issue one license for multiple sources of content, can issue different licenses for the same content, and can support business models that cleanly separate the interests of copyright holders, content distributors, service provider networks, and others. Of course, this basic DRM architecture inherits all the security, privacy, and performance problems of general cross-domain access management. Several authors have already noted the unique security and privacy benefits that Digital Credentials could bring to DRM; see, for instance, <http://www.w3.org/2000/12/drm-ws/pp/hp-poorvi2.html>.

By building DRM on top of CMP, consumers can control and limit the correlations that content distributors can establish about their consuming habits and identity, and content distributors can protect their intellectual properties more securely without infringing fair use rights. Let us walk through a simple CMP-based DRM scenario:

- Bob visits MusicPortal, an Internet portal where the latest album of his favorite band is available via download. MusicPortal groups content distributors together so that customers can buy their music from a single point-of-sale. The portal offers various subscription packages, such as monthly fees for unlimited downloads, prepaid number of downloads, and so on. It also offers Web site personalization and can make recommendations to Bob by keeping track of his musical preferences.
- Bob chooses to purchase a subscription which entitles him to limited number of download every month for a fixed monthly fee. He pays for the subscription with his credit card in a special section of the portal. To protect his privacy, his subscription is delivered in the form of a Digital Credential. This ensures that the subscription cannot be forged, while at the same time the portal will not be able to trace which credit card was used to buy the subscription.
- After making his music selection and the usage rights he wishes to acquire, Bob goes to the checkout section of the portal. To acquire the rights on a specific album Bob presents his subscription to the portal. The portal processes the payment through the clearing house and in exchange emits a digital license describing the rights and privileges associated with the music file. The music file is encrypted specifically for Bob using an encryption key that can be found in the digital license. The digital license is packaged into a Digital Credential as well, to provide lending protection and possibly other protections.
- To play the music, Bob needs a player that understands and enforces the license. Bob's player can permit him to copy the file from one player to another, so that he can play the file from many places. A lending disincentive placed in the licenses (as the credit card information that Bob used to purchase the subscription) would strongly discourage Bob from copying the music to his friends even if he could bypass the hardware protections of his player.
- For extra security, all subscriptions and digital licenses could be managed using a simple 8-bit smartcard, by off-loading all expensive computations and storage to the user's PC, a laptop, or PDA, while preserving all the smartcard's security benefits. Multiple license issuers could all ride along on the security of the same card, without needing to trust each other.

4. Closing Remarks

Currently, the visible battle over user identities is between organizations. The interests of individuals are not seriously taken into consideration by businesses; individuals can only rely on

government legislation and on themselves to reduce the power of organizations to profile them across domains. With ever-increasing advances in data storage, communication, processing, and analysis, both of these are rapidly losing their effectiveness. While some individuals may continue to provide organizations with even more polluted information, others may avoid them altogether. This power struggle between organizations and individuals is in nobody's best interests.

The key to getting everyone on the same side of the table is to adopt identity and access management technologies that give the owners of identity information direct control over how their profile information can be used by others. The notion of ownership of identity information is not always easy to define and capture, however; while often the data subject must be considered to be the legitimate owner of personal information, there are numerous instances where ownership legitimately resides in the hands of one or more organizations, possibly jointly with the data subject. The CMP architecture proposed in this paper has been designed to take this into account; it allows identity-related information to be securely co-managed by both data subjects and organizations, in a manner that simultaneously protects the privacy interests of data subjects and the business and liability interests of organizations.

5. References

- [1] "An In-Depth Analysis of the Liberty Alliance Architecture," Credentica whitepaper, April 2004. <http://www.credentica.com/technology/LA.pdf>
- [2] "A Technical Overview of Digital Credentials," by S. Brands, International Journal on Information Security, 2004 (to appear). <http://www.credentica.com/technology/overview.pdf>
- [3] "Rethinking Public Key Infrastructures and Digital Certificates; Building in Privacy," by S. Brands, 315 pages, August 2000, MIT Press, ISBN 0262-02491-8. With a foreword by prof. Ronald. L. Rivest. <http://www.credentica.com/technology/book.html>



CREDENTIALICA

Non-Intrusive Identity Management

Dr. Stefan Brands

McGill School of Computer Science & Credentica

April 12, 2004

Presented to: 3rd Annual PKI R&D Workshop, Gaithersburg

Personal credentials (1)

credentica.com



- **Adjunct professor at McGill Univ. (Comp. Sc.)**
 - Co-supervising master's and PhD students
 - Member of SSHRC privacy project "On The Identity Trail"
- **Professional cryptographer since 1992**
 - Secure electronic authentication
 - Cross-domain access control & identity management
 - Privacy-enhancing/preserving technologies
 - Electronic payments
- **Member of External Advisory Board of Privacy Commissioner of Canada**
- **Member of CSIS Authentication Working Group**
 - Co-chaired by Lawrence Lessig and Craig Mundie

Personal credentials (2)

credentica.com



- **Author of MIT Press book “Rethinking Public Key Infrastructures; Building in privacy”**
 - Updated version of dissertation (1992-1996 ... 1999)
 - Thesis advisor: prof. Adi Shamir of Weizmann Inst. (**RSA**)
 - Foreword: prof. Ronald L. Rivest of MIT (**RSA**)
 - *“an important landmark in the evolution of privacy-enhancing technology”*
- **Lots of “real-world” experience with privacy & secure authentication technologies:**
 - Consultancy
 - Principal protocol designer of EU-piloted e-purse system
 - Senior cryptographer at privacy-technology companies
 - DigiCash (1996-1998)
 - Zero-Knowledge Systems (2000-2001)
 - Currently with Credentica (2002-...) Funded by Nokia

Content

credentica.com



- **Part I**
 - PKI & digital identity management
- **Part II**
 - Digital Credentials
- **Appendix A**
 - PKI & cross-domain access control (details)



Part I
**PKI & digital identity
management**

PKI – historical perspective

credentica.com



- **1976: Invention of public-key cryptography:**
 - Setting: Message encryption over open network
 - Sender encrypts message with **public key** of recipient
 - To prevent man-in-the-middle attack: rely on on-line database specifying name–public key bindings
- **1978: Kohnfelder (bachelor’s thesis):**
 - Database: bottleneck & vulnerable to attacks
 - **Identity certificates** proposed to address this problem
- **1990’s: X.509 identity certificates provide:**
 - **Confidentiality** of data in transit (through encryption)
 - **User authentication** (ensures messages are encrypted under right public key & prevents man-in-the-middle attack)
 - **Data integrity** (prevent tampering with data in transit)
 - **Non-repudiation** (proof of sender’s identity)

PKI – what is it good for?

credentica.com



- **Message encryption**
 - SSL encryption, secure e-mail
- **Message signing**
 - Document signing, code signing, notarizing, ...
- **“Inescapable” identity**
 - *“Your digital passport to the information highway”*
 - SSL authentication, ...
- **Single-domain access control**
 - Approach (by stretching PKI ...):
 - Access requestor must show identity certificate
 - Certificate = authenticated pointer to back-end database entries
 - Access provider retrieves data for **authorization** decision
 - VPN access, organizational SSO, ...

PKI – what is it NOT good for?

credentica.com



- **Cross-organizational access control**
 - Problems certificates were supposed to address are back (with a vengeance!)
 - Bottleneck of real-time database consulting
 - **Online** database vulnerabilities
 - Need additional security safeguards
 - Cloning of access rights
 - Lending of access rights
 - Privacy problems
 - Like credit card infrastructure on steroids
 - For organizations towards central parties
 - For individuals towards organizations & central parties
- **Cross-organizational identity management**
 - Access control + information sharing/linking/reuse

Digital identity management (1)

credentica.com



- **Network identity**
 - Collection of information relating to an individual
 - Created and managed as single unit in a network
 - Stored in electronic form
- **Situation today:**
 - Individuals have many fragmented network identities
 - Aggregating network identities is increasingly easy
 - But: which network identities pertain to **same** individual?
 - If linkage is wrong, aggregate has less value (**data pollution**)
- **Causes of pollution**
 - Unintentional (different name spellings, ...)
 - Intentional (avoidance, theft, lending, copying, forgery, insider help, ...)

Digital identity management (2)

credentica.com



- **To derive value: network identities must be accompanied by unique cross-domain identifiers**
- **Simple approach:**
 - Request / capture identifier when collecting data
 - Employee ID, static IP address, health insurance number, SSN, credit card number, passport ID, biometric, **X.509 certificate**, ...
 - Aggregate on basis of cross-domain identifier
 - Physical aggregation or logical/virtual aggregation
- **Need to consider different settings**
 - Intra-organizational: Enterprise identity management, ...
 - Extended organizational: extended enterprise (SCM), ...
 - Cross-organizational: E-health, E-government, Critical Information Infrastructures, E-commerce, ...

Where does this approach work fine?

credentica.com



- **Single-domain intra-organizational**
 - Employees have low/no privacy expectations/rights
 - Trust dynamics very simple
 - No scalability issues
 - Can use traditional security tools (“silo protection”)
 - Door guards, intrusion detection, firewalls, anti-virus, ...
- **Multi-domain intra-organizational (branches, ...)**
 - Trust dynamics minimally complicated
 - Low privacy expectations
 - Possibly some scalability issues
 - Traditional security tools still work
- **Large organization with satellite organizations**
 - Authentication relation already unbalanced


Where does it break down?

credentica.com



- **“Balanced” business-to-business (B2B)**
 - Collaborative enterprise applications
- **Government-to-business (G2B)**
- **Business-to-consumer (B2C)**
- **Government-to-citizen (G2C)**
 - Implications for stability of democracy
- **Consumer-to-consumer applications**
 - Online gaming, instant messaging, ...
- **Peer-to-peer transactions**
- ...

Concern over “personal privacy” and “keeping personal information private.”



Privacy & personal information

credentica.com



- **Privacy:** *“The right of individuals, groups, and organizations to determine for themselves when, how, and to what extent information about them is communicated to others.”*
- **Different manifestations for:**
 - Individuals (ROI hard to quantify)
 - Companies (competitive intelligence, liability)
 - Critical Information Infrastructures (monitoring)
- **Personal information:** *“information about a data subject whose identity can reasonably be ascertained from the information”*
- **Network identity = personal information !**
 - Unless **NO PARTY** other than the data subject can determine who is behind a network identity
- **Data protection legislation: organizations must protect personal information**

Fair Information Principles (FIPs)

credentica.com

OECD FIPs:

1. Collection Limitation
2. Data Quality
3. Purpose Specification
4. Use Limitation
5. Security safeguards (incl. confidentiality)
6. Openness
7. Individual Participation
8. Accountability

Technology can address security needs **without** addressing privacy, but this may **introduce grave new security** concerns!

Security safeguards deal mainly with **unauthorized outsiders**, but most privacy threats come from **insiders**

PKI “quick fixes” that do not work (1)

credentica.com



- **Identity certificates specifying a “pseudonym” or a “role” instead of a real name:**
 - Does not address privacy problems (tracing on basis of **public keys/CA signatures** in certificates!)
 - Weakens security (accountability, fraud containment, ...)
 - All other problems remain
- **X.509 attribute certificates**
 - Addresses **only** availability problem
 - Attribute certificates must be linked to (and sent along with) base identity certificate to **prevent pooling of privileges**
 - All other problems **worsen:**
 - **More** privacy-invasive (attributes within certificate known to CA & disclosed when showing certificate)
 - No security mechanisms to prevent discarding, updating-prevention, lending, and cloning (easier when in database!)
 - Must manage and revoke an abundance of certificates

PKI “quick fixes” that do not work (2)

credentica.com



- **Different CA & certificate per domain:**
 - False sense of privacy:
 - Like using SSNs and credit card numbers for all actions
 - Privacy worse due to fully electronic nature
 - Creates “islands” that cannot communicate
 - Greatly reduces functionality for all participants
 - Inform sharing no longer possible
 - SSO goes out the window
 - Only way to link is through bridging CAs
 - This violates what we were trying to achieve
 - Inefficient for client
 - Multiple certificate management
 - Smartcard can hardly handle a single certificate
 - Security limitations
 - no cross-domain revocation

What about federated identity management?

credentica.com



- **Centralize authentication power from different domains into a central domain**
 - Maps cross-domain context back to single-domain context
 - Apply single-domain authentication techniques
 - Password-only, Kerberos, PKI, biometrics, ...
- **Leave authorization decisions at original domains**
- **Liberty Alliance:**
 - Circle of trust: “SPs” using a central “IdP”
 - Counter-movement to Microsoft’s Passport
 - IdP (Microsoft) collected user data, not SPs themselves
 - Allows many circles of trust
 - Allows any single-domain authentication technique
 - User can be known under different pseudonym at each SP

What is it good for?

credentica.com



- **IdP is like a Visa in its “circle of merchants”**
 - Can track, trace, and link all interactions in real time
 - Can impersonate users across circle of trust
 - Can deny access to users across entire circle of trust
 - Appealing target for hackers, insiders, DOS attacks
- **OK if legacy system mirrors this power relation**
 - Intra-organizational identity management
 - Multiple branches
 - Affiliates
 - “Unbalanced” B2B
 - IdP is powerful institution with pre-existing powers over SPs
- **No good in general cross-organizational contexts**
 - User concerns
 - SP concerns

Needed: new authentication primitives!

credentica.com



- **Traditional authentication primitives meet only limited requirements:**
 - Security against unauthorized outsiders
 - Efficiency
 - Scalability
- **Cross-domain access brings NEW requirements**
 - Complicated trust dynamics
 - Vulnerabilities due to (real-time) reliance on central parties
 - Denial of service, hackers, insiders
 - Security against dishonest insiders
 - Perimeter security techniques of little use
 - All security must be tied to the information itself
 - Privacy: Control over who can learn what information



Part II

Digital Credentials

Digital Credentials

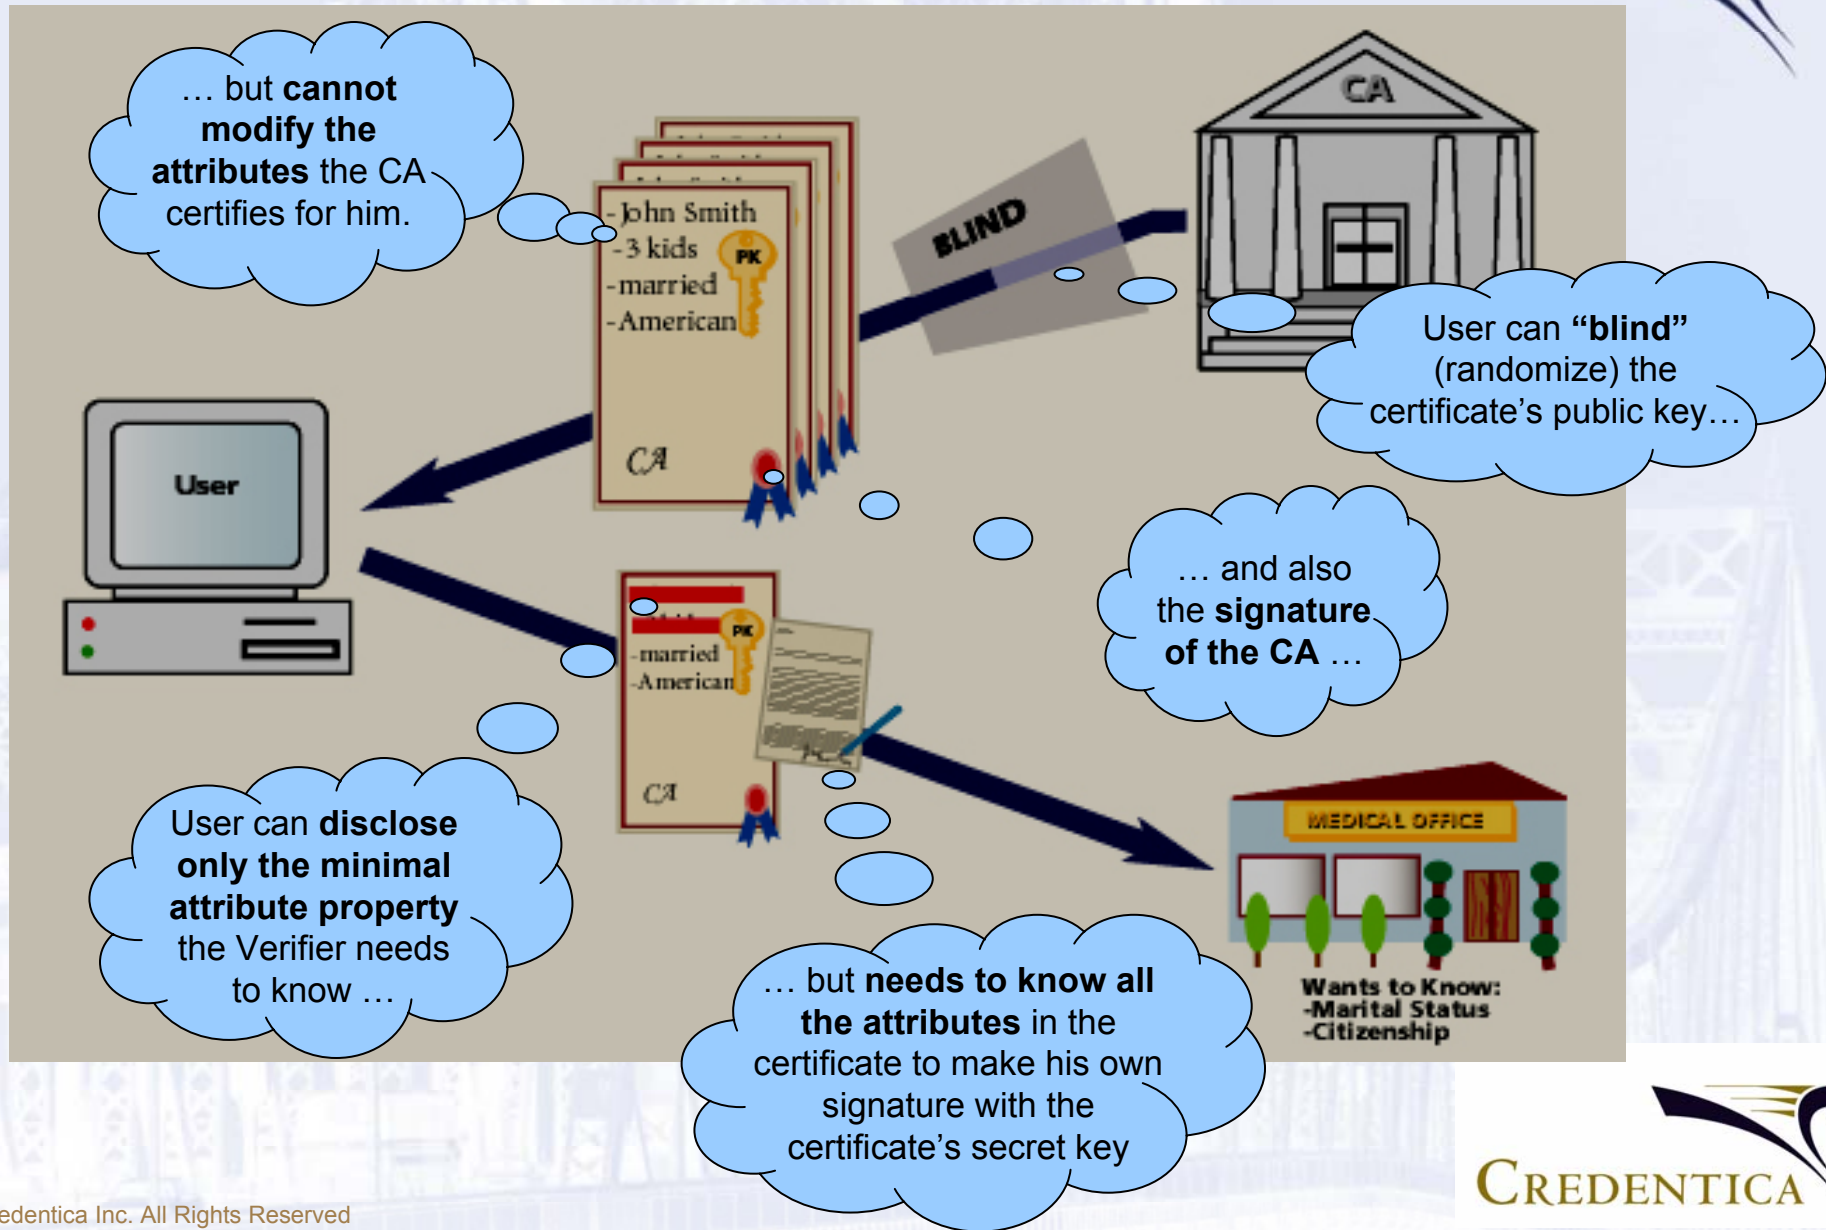
credentica.com



- **Based on 20 years of academic research**
 - Dozens of reputable academics, starting with groundbreaking visionary work by Chaum
- **All the strengths of digital certificates, but much more powerful**
 - One certificate can contain many arbitrary attributes
 - Separate privacy & security sliders
 - **Unlinkable** certificate issuing and showing
 - User can **selectively disclose** attribute properties
 - Identification, anonymity, pseudonymity, “in between”
 - Program can **selectively hide** disclosed attributes
 - Unique security features
 - Efficient smartcard implementation

Digital Credentials in action

credentica.com



Properties of Digital Credentials (1)

credentica.com



- **Fully adaptable levels of privacy:**
 - Anonymous, pseudonymous, role-based access, and anything “in-between”
 - Principle of least authority; **selective/minimal disclosure**
 - Reverse authentication: data does **not** meet conditions
 - **Recertification and updating:** present Digital Credential without revealing current attribute values
 - Dossier-resistance: leave no or partial non-repudiable transaction evidence to verifier
 - Credential verifier can **selectively hide** data before passing on digital evidence 3rd party
 - Credential Authorities can be prevented from learning the attributes that they certify
 - Smartcard cannot leak sensitive data to outside world

Properties of Digital Credentials (2)

credentica.com



- **Security protections:**
 - No **pooling** of privileges (multiple Digital Credentials can be shown to contain same identifier without disclosing it)
 - **Lending** protection: Embed client-confidential data into Digital Credential (legitimate owner need never disclose it)
 - **Discarding** protection: Lump negative data in base Digital Credential (e.g., drunk driving mark into driver's license)
 - **Limited-show** credentials: Embedded identifier (or value) exposed if and only if Credential shown too many times
 - Audit capability:
 - Digital audit trails & receipts facilitate dispute resolution
 - Non-identified audit trail cannot be disavowed by originator
 - Self-signed fraud confessions for lending and reuse

Properties of Digital Credentials (3)

credentica.com



- **Smartcard Implementations:**

- Manage billions of Credentials using 8-bit smart-card chip (off-load storage and computational burden to user device)
- Application provider can arbitrarily minimize level of trust placed in smartcard (through application software)
- Secure multi-application smartcards:
 - Different application providers can share same secret key
 - Digital Credentials have uncorrelated secret keys (unknown even to card supplier) and can be revoked separately
 - Different applications using same smartcard are fire-walled through user software (not card software!)
 - Leakage of a card's key does not allow fraud beyond the security functionality the card was supposed to add

Properties of Digital Credentials (4)

credentica.com



- **Managed services:**
 - Credential Authorities certify sensitive information without being able to learn the data
 - Revocation Authorities can validate certificates without being able to identify the clients of organizations
 - Role of tamper-resistant smartcard can be outsourced
- **Peer-to-peer support:**
 - Individuals can store and manage their own credentials
 - Unauthorized users cannot modify, discard, lend, pool, or prevent the updating of information they hold
 - Distribute all back-end database entries to data subjects
 - Multi-purpose and multi-application certificates

Issuing protocol (example)

USER

CERTIFICATE ISSUER

$$h := g_1^{x_1} g_2^{x_2} g_3^{x_3}$$

$$w_0 \in_{\mathcal{R}} \mathbb{Z}_q$$

$$a_0 := g_0^{w_0}$$

$$\xleftarrow{a_0}$$

$$\alpha_1 \in_{\mathcal{R}} \mathbb{Z}_q^*$$

$$\alpha_2, \alpha_3, w_1, w_2, w_3 \in_{\mathcal{R}} \mathbb{Z}_q$$

$$h' := (hh_0)^{\alpha_1}$$

$$a := (h')^{-w_1} g_1^{w_2} g_3^{w_3}$$

$$c'_0 := \mathcal{H}(h', a, a_0 g_0^{\alpha_2} (hh_0)^{\alpha_3})$$

$$c_0 := c'_0 - \alpha_2 \pmod q$$

$$\xrightarrow{c_0}$$

$$r_0 := (w_0 - c_0)(x_0 + \sum_{i=1}^3 x_i y_i)^{-1} \pmod q$$

$$\xleftarrow{r_0}$$

$$g_0^{c_0} (hh_0)^{r_0} \stackrel{?}{=} a_0$$

$$r'_0 := (r_0 + \alpha_3) \alpha_1^{-1} \pmod q$$

Showing protocol (example)

credentica.com



USER

$$c := \mathcal{H}(h', a, \text{spec})$$

$$r_1 := -c\alpha_1^{-1} + w_1 \text{ mod } q$$

$$r_2 := -cx_1 + w_2 \text{ mod } q$$

$$r_3 := -cx_3 + w_3 \text{ mod } q$$

VERIFIER

$$(h', a), (c'_0, r'_0), x_2, (r_1, r_2, r_3)$$

$$c'_0 \stackrel{?}{=} \mathcal{H}(h', a, g_0^{c'_0} (h')^{r'_0})$$

$$c := \mathcal{H}(h', a, \text{spec})$$

$$(h')^{r_1} a \stackrel{?}{=} g_1^{r_2} g_2^{-x_2 c} g_3^{r_3} h_0^{-c}$$

Details: “Rethinking PKI; building in privacy”

credentica.com

“an important landmark”

Dr. Ronald L. Rivest (Webster Professor of Electrical Engineering and Computer Science at MIT), August 2000

“minimizing the risks of all the interested actors”

Electronic Privacy Information Center & Privacy International, 2001

“a superior alternative to conventional approaches to PKI”

Dr. Roger Clarke (consultant in the management of information and information technology), 2001

“security without sacrificing privacy”


Dr. Hal Abelson (Professor at the Artificial Intelligence Laboratory, MIT), August 2000

“the state of the art”

Dr. A. Michael Froomkin
(Professor of Law, University of Miami), August 2000

“digital certificates without giving so much power to the system owner”

Former Chief Privacy Counselor to the Clinton Administration, Dr. Peter Swire, April 2001





Appendix A

PKI & cross-domain access control (details)

PKI & access control: problems (1)

credentica.com



- **Unscalable beyond pre-established domains:**
 - Access provider relies on the availability, correctness, and timeliness of authorization data
- **Poor security:**
 - Access right cloning and lending: no cryptographic protection
 - Misuse of online databases by hackers and insiders
 - Vulnerable to denial-of-service attacks:
 - Strong reliance on real-time availability of online databases
 - Online certificate status validation
 - Increases risk of identity theft:
 - Inescapable system-wide identification
 - Strong reliance on central databases

PKI & access control: problems (2)

credentica.com



- **Not suitable for use with smartcards:**
 - Cannot use low-cost smartcards:
 - Storage problem
 - Need crypto co-processor for exponentiations
 - Elliptic-Curve cryptography is only partial solution
 - Application provider **must place very strong trust** in parties involved in smartcard manufacturing, masking, initialization, application loading, and personalization.
Attacks:
 - Overt or covert leakage of secrets and other confidential data
 - Uniqueness, randomness, and secrecy of secret keys??
 - Fake-terminal attacks
 - Selective “failure” attacks based on dynamic inputs
 - Problems worsen for multi-application smartcards

PKI & access control: problems (3)

credentica.com



- **Managed services are intrusive:**
 - Online Certificate Status Providers able to learn competitive/sensitive data in real time:
 - Identities of access requestors (and access providers)
 - Peak hours
 - Typically: nature of the transaction
 - Possibly: transaction details
 - Certificate Authorities must know the identity and any other attributes that go into the certificates they issue
 - Online Certificate Status Providers & Certificate Authorities & on-line database maintainers can **disrupt operations** on the basis of transaction-specific knowledge **in real time**

PKI & access control: problems (4)

credentica.com



- **Privacy-invasive (inescapable systemic identification rooted into infrastructure):**
 - Public keys = strongly authenticated “super-SSNs”:
 - **Globally unique** identification numbers
 - **Inescapably travel along** with each and every action taken
 - Obtained by access provider **& third parties** (providers of authorization databases & online certificate status verifiers)
 - Always leave behind **undeniable digital evidence** of the requestor’s identity (due to digital signing of nonces)
 - Problems with data protection legislation, unbridled use of PKI may be unconstitutional
 - Access providers & 3rd parties **cannot prevent** receiving identifiable data

Role Sharing in Password-Enabled PKI

Xunhua Wang[†], Samuel Redwine
Commonwealth Information Security Center &
Department of Computer Science
James Madison University
Harrisonburg, VA 22807 USA
{wangxx,redwinst}@jmu.edu

Abstract

Password-enabled PKI schemes simplify the management of end users' private keys by storing them in password-protected form on a centralized on-line server. Under such schemes an end user needs only remember his password and can access his private key from anywhere the centralized server is available. Existing password-enabled PKI schemes are based on the single-user model where a private key is owned by one user. In this article, we present mechanisms to support role sharing in password-enabled PKI. In our schemes, using passwords only, a group of users share the privileges of a role through sharing the private key of that role. We first develop a hybrid password-enabled PKI scheme, which supports both easy password change and misuse monitoring. Then, based on this hybrid and existing password-enabled PKI schemes, we give password-enabled role sharing schemes for both threshold access structures that require a threshold number of these users to execute the shared role and more general access structures that allow more flexible role sharing policies.

Keywords: Role sharing, Password-enabled PKI, PAKE

1 Introduction

In a password-enabled PKI scheme [22, 17], the private key of an end user is not stored on a smart-card or on the user's laptop. Instead, it is protected by a password chosen by the user and stored on a centralized online server. Compared to the conventional smartcard-based PKI approach, password-enabled PKI is a lightweight solution and enjoys high usability: no smartcard reader is required; an end user needs only remember his password and can roam anywhere the centralized server is available.

There are two different approaches for password-enabled PKI, *virtual soft token* [20, 17] and *virtual smartcard* [22]. In the virtual soft token PKI [20, 17], a password is used to encrypt the private key of a public/private key pair and the encrypted private key is stored on a centralized server. With his password, a user can remotely authenticate himself to the server, establish an authenticated and cryptographically strong session key (thus, a secure connection) with the server, download the encrypted private key via the secure channel, decrypt it and use the private key as in the conventional PKI activities. The first step of this approach authenticates a user before he can download a password-encrypted private key and the second step establishes a session key to protect the subsequent downloading of the password-encrypted private key from the off-line

[†]A part of this work has been supported by a Cisco CIAG grant.

dictionary attack [19]. These two steps can be accomplished by a *password-authenticated key exchange (PAKE)* protocol [2, 16, 30].

In the virtual smartcard PKI [22], an end user's private key is split into two parts, a *human memorizable password* and a *server component*. The end user holds the password and the server component is stored on a server. Like in the virtual soft token, to use his private key, a user of the virtual smartcard PKI first runs the PAKE protocol with the server to have mutual authentication and establish a secure channel. Then, the user applies his password to a message (either a message to be digitally signed or a ciphertext to be decrypted) and sends the partial result to the server over the secure channel. The server combines its own partial result, computed from the corresponding server component and the message, with the user's partial result to generate the final result. The major difference between virtual soft token schemes and virtual smartcard schemes is that, in virtual smartcard schemes, every cryptographic operation (such as digital signature and decryption) requires the cooperation of the centralized server while in virtual soft token schemes an end user can do many cryptographic operations as he wants after securely downloading his private key.

The problem. All existing password-enabled PKI schemes [20, 17, 22] are based on the one-user-one-private-key model and are essentially *single user-oriented*. That is, they do not support multiple-users-one-private-key and thus do not support *role sharing*.

A *role*, in the access control community, is defined as a basic semantic unit to describe the authority and responsibility that users of that role assume [24]. A good organization-wide access control decision is often based on roles (for example, president of AOL), instead of any specific individual user, as users may change over time while roles change less frequently. In many role-based access control models [24, 23], a user assigned to a specific role is implicitly granted all the privileges of that role. However, as pointed out in [6], within an organization, the responsibility of a role is not always assumed by any single individual but sometimes is shared among a group of users of that organization. To execute the role privileges, a subgroup of these users are required to agree on the action. In this way, power abuse by a single user or a small coalition of these users can be prevented and the principle of separation of duty can be guaranteed. We consider the case where a public/private key pair, called *role public/private key pair*, is affiliated with the role. The role public key is used by external users to encrypt messages intended for this role or verify messages digitally signed by the role private key. For users external to the role, what they see is the role itself and the users assigned to this role are invisible to them. Possibly, when collectively executing the role privileges, the users sharing a role do not necessarily trust each other.

Based on the above observation, in this article, we explore password-enabled PKI schemes to support role sharing, which are called *password-enabled role sharing PKI*.

Our contribution. We first propose a new password-enabled PKI scheme, called *hybrid password-enabled PKI*, which is later extended to support role sharing. Compared to existing password-enabled PKI schemes [20, 22, 17], this hybrid scheme allows server administrators to perform instant revocation of a user's public key and to monitor user PKI activities for misuse detection and, at the same time, it supports user password change very well. (Previous password-enabled PKI schemes support only one of these two features.) Then, we propose *password-enabled role sharing PKI* schemes for both threshold access structure and general access structure. In the scheme for threshold access structure, which is called *threshold password-enabled role sharing PKI*, a group of (say n) users are assigned to a role (and thus share the role private key), each with his favorite password and nothing else, and a threshold (say, t , $t \leq n$) of them are required to cooperate to execute the role privileges without reconstructing the shared role private key at any single location. Like the traditional single user-oriented password-enabled PKI schemes, our architecture adopts a central server, where password-protected credentials are stored. A subset of users fewer than t , together with the centralized server, will not be able to use the role private key directly. In

Table 1: Comparison of our work with previous research

Single User-oriented Password-enabled PKI		Password-enabled Role Sharing PKI	
Name	Property	Name	Property
Virtual soft token	easy password change	Threshold virtual soft token [‡]	easy
		Type-1 password-enabled role sharing PKI for general access structure [‡]	password change
Virtual smartcard	misuse detection		
Hybrid password-enabled PKI [‡]	easy password change & misuse detection	Threshold hybrid password-enabled PKI [‡]	easy password change & misuse detection
		Type-2 password-enabled role sharing PKI for general access structure [‡]	

this article we also propose password-enabled role sharing schemes to support more general access structure, in which more flexible role sharing policies are allowed.

It should be noted that our role-sharing schemes are different from the voting-based role sharing approach, where a *fully trusted* centralized server checks the votes from a subgroup of users and, if a certain condition is met, executes the role privilege. In our schemes, the centralized server is not fully trusted and the server itself alone cannot directly execute the role. Thus, neither the central server administrator nor an attacker who has successfully compromised the server can assume the role directly. Table 1 gives comparison between this work (marked with [‡]) and previous research. The first two columns of table 1 give the single user-oriented password-enabled PKI schemes, including the hybrid password-enabled PKI proposed in this paper, and the last two columns of the table list their corresponding extensions for role sharing.

The remainder of this article is organized as follows. Section 2 gives the related work. Section 3 presents a hybrid password-enabled PKI scheme. Section 4 discusses some principles for designing role sharing password-enabled PKI schemes. Section 5 presents our role sharing password-enabled PKI schemes for threshold access structure and 6 gives our role sharing password-enabled PKI schemes for general access structures. In Section 7 we discuss some operational and performance issues. Concluding remarks are given in Section 8.

2 Related Work

Desmedt [6] first proposed the concept of group-oriented cryptography to allow a threshold number of users sharing a group private key. In all the threshold cryptography schemes, including those threshold RSA [10, 8, 9, 21, 26, 14] and threshold DSS [12, 13, 18] schemes, each user of the role is assigned one or more *long* random secret shares of the role private key. Since most human being are not good at memorizing long random secrets and smart-cards have not been widely used yet, so far these threshold cryptography schemes have only been used in machine-oriented applications [31, 1, 29], not people-oriented systems. In contrast, the schemes explored in this article are password-based and thus, people-oriented.

Ganesan [11] first introduced passwords into the 2-out-of-2 threshold RSA [5] and used it to enhance the Kerberos system. We notice that this enhancement, like [22], is still single-user oriented and does not support role sharing.

Using a 2-out-of-2 threshold RSA scheme Boneh et al. [4] proposed an architecture for fast public key revocation. In their architecture is a semi-trusted mediator (SEM) who can monitor a user's

PKI activities. Compared to this scheme, our hybrid password-enabled PKI scheme (presented in Section 3) is password-enabled and thus enjoys better usability.

3 A New Password-Enabled PKI Scheme

The virtual soft token PKI scheme proposed in [20] allows its users to change their password easily. However, the administrators of its centralized server cannot monitor users' PKI activities as a user can perform many PKI operations after downloading his private key. On the other hand, the virtual smartcard PKI scheme proposed in [22] allows the administrators of the centralized server to monitor user PKI activities and supports instant public key revocation. However, as observed in [28], user password change is not supported very well as it is computation intensive.

In this section, we propose a hybrid password-enabled PKI scheme that supports both PKI activity monitoring and simple password change. The essential idea behind the hybrid password-enabled PKI scheme is that an additive 2-out-of-2 secret sharing is performed on the user's private key first and one of the two resulting shares, called the *server component*, is assigned to the centralized server. The other share, called the *client component*, is assigned to the user and is encrypted with the user's password and stored on the centralized server. When the user needs to use his private key, he securely downloads the password-encrypted key share, as done in virtual soft token, decrypts the key share and uses it to compute a partial result. To get the final result, the user needs the cooperation of the centralized server, which uses the server component as in the virtual smartcard scheme. Thus, this hybrid password-enabled PKI scheme is similar to the virtual smartcard scheme [22] in that the centralized server is also assigned a component of the user's private key; on the other hand, it is also similar to the virtual soft token [20] in that a user needs to download a password-encrypted credential to perform the client-side computation, which makes password change simpler.

Below we give the details of the RSA-type hybrid password-enabled PKI scheme. The same idea can be used to build DSA-type hybrid password-enabled PKI scheme but it is more complicated [18].

RSA-type hybrid password-enabled PKI Assume that Alice is a user of the hybrid password-enabled PKI scheme and her RSA public key is (N, e) , where $N = p \times q$, p and q are two primes. d is Alice's corresponding private key.

- **Component generation.** In our hybrid password-enabled PKI scheme, the centralized server picks a random r , $1 \leq r \leq \phi(N)$ where $\phi(N) = (p - 1) \times (q - 1)$, and computes $r' = d - r \bmod \phi(N)$. r is the server component and r' is the client component. Alice picks her favorite password \hat{p} and uses it to encrypt r' . The password-encrypted result, $y = E_{\hat{p}}(r')$, is stored on the server. For Alice, the centralized server also stores a password verification data which is a value derived from \hat{p} and is used by the server to run a PAKE protocol with Alice.
- **Private key use.** Armed with her password, \hat{p} , Alice runs a PAKE protocol with the centralized server and establishes a secure channel. She then securely downloads y and decrypts at the client side to recover r' . To use her private key to perform a cryptographic operation on a message m , Alice first applies r' to m to get a partial result $c_1 = m^{r'} \bmod N$. c_1 is sent to the centralized server via the secure channel and the server applies its r to m to get $c_2 = m^r \bmod N$. The final result is $c_1 \times c_2 \bmod N$, which is equivalent to applying Alice's private key on message m .

In the above process, the centralized server is required to participate in the computation, which allows the centralized server administrator to monitor users' PKI activities and do instant public key revocation. On the other hand, Alice can change her password \hat{p} to another password \bar{p} by downloading y , recovering r' , computing $y' = E_{\bar{p}}(r)$ and sending it to the centralized server. None of these steps is computation intensive and can be simply performed.

It is worth mentioning that in this hybrid scheme, every cryptographic operation related to the private key requires interactions with the centralized server. In contrast, with a virtual soft token, a user can load his private key onto the laptop and work offline, decrypting emails and signing new messages with no further interactions with the server.

4 Design Principles for Password-enabled Role Sharing PKI

In the remainder of this paper, PU_{role} and PK_{role} are used to denote the role public key and the role private key respectively. n is the size of the group of users to share the role and we use $\mathcal{P} = \{U_1, U_2, \dots, U_n\}$ to denote the set of the users. An *authorized subset* is defined as a subset of \mathcal{P} whose users are allowed to collectively execute the role privileges and an *access structure*, Γ , for the role is the set of authorized subsets [27, pages 331].

4.1 Centralized server

Besides the users to share a role, in our architecture, there is a centralized on-line server, as in the traditional single user-oriented password-enabled PKI schemes [20, 17, 22]. It is this on-line server that makes password-enabling possible. On the other hand, this on-line server is not fully trusted in the sense that role-related credentials are *not* stored in the clear on it, but protected by passwords, and the private credentials are never exposed on the server. This distinction differentiates both the traditional password-enabled PKI and our schemes from the voting-based approach where the server is fully trusted.

For each user sharing a role, after he picks a password, the centralized server also stores the corresponding password verification data (PVD) for that user.

4.2 Design principles

There are several design principles for our password-enabled role sharing schemes. Some are straightforward while others are not.

1. The role public/private key pair does not change as often as the users assigned to the role. This is the rationale for role-based access control and is also true in our role sharing password-enabled PKI schemes.
2. A user revoked from a role should *not* know the shared role private key. Nor do a small coalition of users who have been revoked from the same role and who are not in the access structure anymore. Obviously, virtual soft token [20] does not meet this principle.
3. Users sharing a role possess passwords only and nothing more. All operations related to the role need the explicit permission of users from an authorized subset.
4. No full trust is placed on the centralized server. The server should not know the role private key. Thus, its administrators or a hacker who has compromised the server cannot execute

the role privilege in a simple way. On the other hand, using what's stored on the server, the server administrators can mount off-line dictionary attacks. This characteristics is common to all password-based schemes and can be mitigated using multiple servers [28]. We notice that *not* all passwords are vulnerable to off-line dictionary attacks. Moreover, compared to the traditional single user-oriented password-enabled PKI schemes, in our role sharing schemes, it is harder for a malicious server administrator or a hacker who has taken control the server to mount off-line dictionary attacks as multiple, instead of a single, passwords are involved.

5. Both threshold access structure and general access structure should be supported. In a threshold access structure, any subset of size not less than the threshold is an authorized subset and this access structure is commonly used. On the other hand, threshold access structure is not always applicable and sometimes more general access structure is used.

5 Threshold Password-enabled PKI

In this section, we shall present *threshold password-enabled PKI* schemes. In the following discussion, t , $t \leq n$, is the threshold. We first discuss how to add role sharing support to the virtual soft token scheme [20]. We then extend the hybrid password-enabled PKI proposed in Section 3 to support role sharing.

5.1 Threshold virtual soft token

Threshold virtual soft token is the role sharing extension of the virtual soft token scheme [20]. Threshold cryptography schemes [7, 12, 26] are used to for this purpose. We have two types of threshold virtual soft tokens, the *threshold virtual RSA soft token* for RSA-type role public/private key pair and the *threshold virtual DSA soft token* for DSA-type role key.

In a threshold virtual RSA (DSA) soft token, a role RSA (DSA) public/private key pair is first generated and then shares of the role private key, PK_{role} , are generated through a (t, n) Shamir secret sharing [25], $(s_1, s_2, \dots, s_n) \xrightarrow{(t,n)} PK_{role} \bmod \phi(N)^*$ where s_i are the shares. Each user U_i , $1 \leq i \leq n$, picks his password, \hat{p}_i , and his corresponding password verification data, PVD_i , is generated and stored on the centralized server. For each user, also stored on the centralized server is y_i , the encryption of s_i by \hat{p}_i (that is, $y_i = E_{\hat{p}_i}(s_i)$).

When the role privilege needs to be executed, depending on the threshold cryptography scheme employed, users' steps vary. In our following discussions we use the threshold RSA given in [26], called Sho00, and the threshold DSA scheme given in [12, 13], called GJKR96.

Threshold virtual RSA soft token To authorize a role-related operation, t or more of the n users are required. Let m be the message to be processed by the role private key. Each participating user first uses his password to run a PAKE protocol with the centralized server and establishes a secure connection; he then securely downloads the password-encrypted key share s_i , decrypts it and computes a partial result as $c_i = m^{2\Delta s_i} \bmod N$ where $\Delta = n!$ [26]; c_i is sent back to the centralized server over the secure channel. After collecting enough partial results, the centralized server combines them into the final result. The Sho00 threshold RSA is non-interactive and thus, in the role execution, users do not need to interact with others.

*For DSA, the modulus is the DSA system parameter q .

Threshold virtual DSA soft token When t or more users want to collectively authorize a role-related operation on message m , each of them first runs a PAKE protocol to log onto the centralized server, securely downloads his y_i and decrypts it as s_i . They then use the GJKR96 threshold DSA to collectively generate a DSA signature on m . The GJKR96 threshold DSA is an interactive scheme while, in our applications, interactions between users are not desirable. Fortunately, we observe that the interactive computation (all the steps until the computation of r [13, pages 70]) of the GJKR96 threshold DSA scheme are message-independent and can be pre-computed. Based on this observation, in our threshold virtual DSA soft token, we can avoid user interactions by performing the message-independent interactive computations in a *partially-protected* store-and-forward way: all the broadcast messages by user U_i are sent to the centralized server in the clear, which will be forwarded to other participating users by the server, and all the *intermediate* private messages of U_i are encrypted by U_i 's password before they are sent to and stored on the server (for future use). These pre-computations need *no* input from users and can be performed, without user U_i 's interventions and notices, after U_i logs into the system.

In GJKR96, b , the number of users required for a threshold DSA signature, is $(2t - 1)$, not t . That is, t should satisfy that $t < \frac{n}{2}$. Therefore, in our threshold virtual DSA soft token scheme, $(2t - 1)$ users are required to collectively execute the shared role.

Both the threshold virtual RSA soft token and threshold virtual DSA soft token allow a user to change his password while keeping his role private key share unchanged. To change his password, U_i uses his old password to run a PAKE protocol with the server, securely downloads the key share protected by the old password, decrypts it, re-encrypts it with his new password, and securely uploads it to the server. To change his password, the user should also notify, via the secure connection, the server of his new PVD.

In the above threshold virtual soft token schemes, although the centralized server is used as a working platform, it is not assigned a share of the role private key and does not contribute to the final result.

5.2 Threshold hybrid password-enabled PKI

In a virtual smartcard scheme [22], the centralized server is also assigned a share of the user's private key and is required to participate in the computation when the user's private key is used. This allows an administrator of the central server to monitor the use of the user's private key and to instantly disable the user's private key if his public key is revoked. (In contrast, the virtual soft token [20] scheme does not offer this monitoring granularity since the private key is recovered and used on the user's machine.)

It is not immediately obvious on how to extend the virtual smartcard scheme given in [22] to support password-enabled role sharing. In a (t, n) Shamir secret sharing scheme [25], to share a secret, at most $(t - 1)$ shares can be passwords. This fact prevents us from simply extending the virtual smartcard scheme for password-enabled role sharing since, ideally, in a password-enabled role sharing scheme, all the n , not just $(t - 1)$, users hold their favorite passwords only and nothing else. One might think to apply the following extension to the virtual smartcard scheme: for each combination $U_{i_1}, U_{i_2}, \dots, U_{i_t}$, where $\{i_1, i_2, \dots, i_t\} \subset \{1, 2, \dots, n\}$, we can compute $d_{\{i_1, i_2, \dots, i_t\}} = d - \hat{p}_{i_1} - \hat{p}_{i_2} - \dots - \hat{p}_{i_t} \bmod \phi(N)$, where \hat{p}_{i_j} is the password of U_{i_j} , $1 \leq j \leq t$, and store $d_{\{i_1, i_2, \dots, i_t\}}$ on the server. In this way, any t users can cooperate with the server to collectively apply the shared role private key on a message. However, this extension has a security flaw: it stores $\binom{n}{t}$ such $d_{\{i_1, i_2, \dots, i_t\}}$ values on the server and in some cases the server will be able to restore the role private key from them, which contradicts with our design principle 4 (see Section 4).

On the other hand, the hybrid password-enabled PKI scheme proposed in Section 3 can be extended to support role sharing for threshold access structure, which allows both monitoring granularity and easy password change. The following details are based on the RSA-type hybrid password-enabled PKI give in Section 3.

- **Component generation.** After the role RSA public/private key pair $(PU_{role}, PK_{role} = d)$ is generated, a random r , $1 \leq r \leq \phi(N)$, is generated and r' is computed as $r' = d - r \bmod \phi(N)$. r is the server component and is stored on the centralized server. A (t, n) Shamir secret sharing is performed on the client component r' , $r' \xrightarrow{(t,n)} (s_1, s_2, \dots, s_n) \bmod \phi(N)$ and s_i is assigned to U_i , $1 \leq i \leq n$. Each user U_i picks his password, \hat{p}_i , and it is used to encrypt s_i into $y_i = E_{\hat{p}_i}(s_i)$. For user U_i , y_i and a password verification data derived from \hat{p}_i are stored on the centralized server.
- **Private key use.** When t or more users agree to apply the role's private key on a message m , each of them uses his \hat{p}_i to run a PAKE protocol with the centralized server and establish a secure channel. He then securely downloads y_i and decrypts at the client side to recover s_i . He then first applies s_i to m and gets a partial result $c_{1i} = m^{s_i} \bmod N$. c_{1i} is sent to the centralized server via the secure channel. The server also applies its r to m to get $c_2 = m^r \bmod N$. After collecting enough partial results, the centralized server combines all partial results, c_{1i} and c_2 into the final result, which is exactly of the role's private key on message m .

In the above process, the centralized server is required to participate, which allows the centralized server administrator to monitor the role's PKI activities and do instant public key revocation. On the other hand, each user can change his password \hat{p}_i to another password \bar{p}_i by downloading y_i , recovering s_i , computing $y'_i = E_{\bar{p}_i}(s_i)$ and sending it to the centralized server. None of these steps is computation intensive and can be simply performed.

6 Password-enabled Role Sharing for General Access Structure

In our real world not all access structures are threshold-based and sometimes more general access structures are used. For example, four users, (U_1, U_2, U_3, U_4) , share a role and $\Gamma = \{\{U_1, U_2\}, \{U_1, U_3, U_4\}\}$ is its access structure. This access structure is not threshold: $\{U_1, U_2\}$ has two members and is allowed to execute the role while $\{U_2, U_3, U_4\}$ is not allowed although its cardinality is 3.

In this section we discuss how to support password-enabled role sharing for general access structure. An access structure is said to be *monotone* if $B \in \Gamma$ and $B \subseteq C \subseteq \mathcal{P}$ implies $C \in \Gamma$ [3]. We are only interested in monotone access structure here.

General access structure-oriented secret sharing — called *generalized secret sharing* — was first studied by Ito et al. [15]. Benaloh and Leichter [3] developed a simpler generalized secret sharing, which is called BL88 in the following discussion. It should be noted that password-enabled role sharing discussed here is more than secret sharing as we do not reconstruct a shared secret, as done in secret sharing schemes, since reconstruction leads to a single point of attack.

In the rest of this section we will give two types of password-enabled role sharing PKI schemes for general access structure. Our discussions are based on the RSA algorithm but can also be applied to DSA.

6.1 Type-1 password-enabled role sharing PKI

Type-1 password-enabled role sharing PKI for general access structure is the extension of the virtual soft token for role sharing.

- Component generation. Given a monotone access structure Γ for a role whose private key is $PK_{role} = d$, we first use the BL88 generalized secret sharing scheme to generate secret shares. Each of the n users will get one or more secret shares. Then, each of them picks his favorite password \hat{p}_i and uses it to encrypt all of his secret shares. The password-encrypted secret shares, together with a password verification data derived from \hat{p}_i , are stored on the centralized server.
- Private key use. When an authorized subset of users want to execute the role privilege on a message m , each of them, U_i , runs a PAKE protocol to log onto the centralized server and establishes a secure connection with it. U_i then securely downloads his secret shares and applies it to m to get a partial result. The partial result is securely sent to the centralized server who combines all partial results into a final result, which is equivalent to applying the role's private key on m .

Using the above $\Gamma = \{\{U_1, U_2\}, \{U_1, U_3, U_4\}\}$ as an example, we have the following shares: d_1 is assigned to U_1 , d_2 is assigned to U_2 , d_3 is assigned to U_3 , d_4 is assigned to U_4 where $d_1 + d_2 = d \bmod \phi(N)$ and $d_1 + d_3 + d_4 = d \bmod N$. When U_1 and U_2 agree to apply the role private key to m , U_1 computes $c_1 = m^{d_1} \bmod N$ and U_2 computes $c_2 = m^{d_2}$. After receiving c_1 and c_2 , the centralized server combines them into the final result as $c = c_1 \times c_2 \bmod N = m^d \bmod N$, which is exactly the role private key on m .

In the above steps, the centralized server does not contribute to the final result and technically the step of combining partial results into the final result can be performed by any users. That is, type-1 password-enabled role sharing PKI does not provide a technical means to monitor role PKI activities on the centralized server.

6.2 Type-2 password-enabled role sharing PKI

Using the same idea of the hybrid password-enabled PKI, type-1 password-enabled role sharing PKI can be modified so that the centralized server is required to contribute for a role privilege execution. In the above component generation stage, instead of sharing d , we can first run a 2-out-of-2 additive secret sharing on d and get d' and d'' , where $d = d' + d'' \bmod \phi(N)$. d' is the server component and is assigned to the centralized server. The client component, d'' , is shared among the n users using the BL88 generalized secret sharing. Then each user uses his password to encrypt the shares assigned to him and stores the password-protected shares on the centralized server. When an authorized subset of users want to execute the role privilege, they compute their partial results. To get the final result, the centralized server is also required to participate and compute its own partial result. Thus, this modified scheme allows the centralized server administrators to monitor role PKI activities and is called *type-2 password-enabled role sharing PKI for general access structure*.

For general access structure, a user is likely to be assigned more than one secret shares and, in deciding which share to use, he needs to know the identities of others users of the authorized subset. This might be undesirable sometimes as it needs coordinations between the participating users. A method for U_i to avoid this interaction is to apply all of his secret shares to m to get more partial results than necessary. When the centralized server combines the partial results, only those necessary partial results will be used.

7 More Discussions

In this section we discuss some performance and operational issues.

7.1 Performance considerations

Compared to the virtual soft token [20] and the virtual smartcard scheme [22], the hybrid password-enabled PKI scheme does not introduce any additional significant computational cost.

7.2 Operational considerations

Password-based versus smartcard-based. Passwords are commonly used for authentication in our daily lives and support user roaming very well. Password-enabled PKI schemes integrate the roaming capability and good usability of passwords into PKI. However, in some application cases, smartcard-based solution might still be preferred due to its high-level security. For these applications, password-enabled PKI can be used as a short-term solution and the migration from password-based to smartcard-based can be made smooth.

In both threshold virtual soft token scheme and threshold hybrid password-enabled PKI scheme, an end user is assigned one or more shares of the role private key and the password-encrypted key shares are stored on the centralized server. This structure makes it easy for password and smartcards to co-exist and makes it easy to migrate from password-based to smartcard-based: users who prefer passwords can still hold their passwords and store their password-encrypted shares on the server; users who like smartcards can download their password-encrypted key shares and feed them to smartcards. This is also true for both type-1 and type-2 password-enabled role sharing PKI schemes.

Recovery from password loss. In a password-based system, a user might inadvertently lose his password to somebody else. For example, a user might use an insecure computer on which a key logging program is installed to harvest passwords. For single user-oriented password-enabled PKI, this might be disastrous. In contrast, the password-enabled role sharing PKI schemes tolerate this type of mistakes to some extent: as long as an attacker does not steal more than $(t - 1)$ passwords, he will not be able to assume the shared role. The recovery from such loss is also straightforward: after a user loses his password, his old key share is disabled and any t other users who share the same role can help him get a new key share.

8 Conclusion

Conventional password-enabled PKI schemes are based on the one-private-key-one-user model and do not support role sharing. In this article we developed schemes to add role sharing to password-enabled PKI schemes. We first presented a hybrid password-enabled PKI scheme, which supports both easy password change and misuse monitoring. Then, we extended our hybrid and existing password-enabled PKI schemes to support role sharing. Our password-enabled role sharing PKI schemes support both threshold access structures and general access structures. Compared to conventional password-enabled PKI schemes, from an end user's perspective, our password-enabled role sharing PKI schemes do not incur additional significant computational cost and also tolerate end user's operational mistakes.

9 Acknowledgement

We wish to thank the anonymous reviewers for pointing us to the related work in [4] and for other useful comments.

References

- [1] B. Barak, A. Herzberg, D. Naor, and E. Shai. The proactive security toolkit and applications. In *Proceedings of the 6th ACM Conference on Computer and Communications Security*, pages 18–27, November 2–4 1999.
- [2] S. Bellare and M. Merritt. Encrypted key exchange: password-based protocols secure against dictionary attacks. In *Proceedings of the 1992 IEEE Computer Society Symposium on Research in Security and Privacy*, pages 72–84, 1992.
- [3] J. Benaloh and J. Leichter. Generalized secret sharing and monotone functions. In S. Goldwasser, editor, *Advances in Cryptology - Crypto '88*, number 403 in *Lecture Notes in Computer Science*, pages 27–36, Berlin, 1988. Springer-Verlag.
- [4] D. Boneh, X. Ding, G. Tsudik, and C. M. Wong. A method for fast revocation of public key certificates and security capabilities. In *Proceedings of the 10th USENIX Security Symposium*, August 13–17 2001.
- [5] C. Boyd. Some applications of multiple key ciphers. In C. G. Gunther, editor, *Advances in Cryptology, Proc. of Eurocrypt '88*, volume 330 of *Lecture Notes in Computer Science*, pages 455–467, Davos, Switzerland, May 1988. Springer-Verlag.
- [6] Y. Desmedt. Society and group oriented cryptography: a new concept. In *Advances in Cryptology, Proc. of Crypto '87*, pages 120–127, August 16–20 1988.
- [7] Y. Desmedt. Some recent research aspects of threshold cryptography. In E. Okamoto, G. Davida, and M. Mambo, editors, *Information Security*, volume 1396 of *Lecture Notes in Computer Science*, pages 158–173, September 1997. URL <http://www.cs.fsu.edu/~desmedt/ISW.pdf>.
- [8] Y. Desmedt, G. Di Crescenzo, and M. Burmester. Multiplicative nonabelian sharing schemes and their application to threshold cryptography. In *Advances in Cryptology — Asiacrypt '94*, pages 21–32, November/December 1994.
- [9] Y. G. Desmedt and Y. Frankel. Homomorphic zero-knowledge threshold schemes over any finite abelian group. *SIAM Journal on Discrete Mathematics*, 7(4):667–679, November 1994.
- [10] Y. Frankel and Y. Desmedt. Parallel reliable threshold multisignature. Tech. Report TR-92-04-02, Dept. of EE & CS, Univ. of Wisconsin-Milwaukee, April 1992. ftp://ftp.cs.uwm.edu/pub/tech_reports/desmedt-rsa-threshold_92.ps.
- [11] R. Ganesan. Yaksha: Augmenting Kerberos with public-key cryptography. In *Proceedings of the ISOC Network and Distributed Systems Security Symposium*, 1995.
- [12] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Robust threshold DSS signatures. In *Advances in Cryptology — Eurocrypt '96*, pages 354–371, May 12–16 1996.

- [13] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Robust threshold DSS signatures. *Information and Computation*, 164(1):54–84, 2001.
- [14] N. Gilboa. Two party RSA key generation. In M. Wiener, editor, *Advances in Cryptology - CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 116–129, Santa Barbara, California, USA, August 1999.
- [15] M. Ito, A. Saito, and T. Nishizeki. Secret sharing schemes realizing general access structures. In *Proc. IEEE Global Telecommunications Conf., Globecom'87*, pages 99–102. IEEE Communications Soc. Press, 1987.
- [16] D. P. Jablon. Strong password-only authenticated key exchange. *ACM SIGCOMM Computer Communication Review*, 26(5):5–26, October 1996.
- [17] T. Kwon. Virtual software tokens - a practical way to secure PKI roaming. In G. Davida, Y. Frankel, and O. Rees, editors, *Proceedings of the Infrastructure Security (InfraSec)*, volume 2437 of *Lecture Notes in Computer Science*, pages 288–302. Springer-Verlag, 2002.
- [18] P. MacKenzie and M. Reiter. Two-party generation of DSA signatures (extended abstract). In J. Kilian, editor, *Advance in Cryptology - EUROCRYPT 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 137–154, Santa Barbara, CA, USA, August 2001. Springer.
- [19] R. Morris and K. Thompson. Password security: a case history. *Communications of the ACM*, 22(11):594–597, November 1979.
- [20] R. Perlman and C. Kaufman. Secure password-based protocol for downloading a private key. In *Proceedings of the ISOC Network and Distributed Systems Security Symposium*, 1999.
- [21] T. Rabin. A simplified approach to threshold and proactive RSA. In *Advances in Cryptology, Proc. of Crypto'98*, pages 89–104, August 23-27 1998.
- [22] R. Sandhu, M. Bellare, and R. Ganesan. Password enabled PKI: Virtual smartcards vs. virtual soft tokens. In *Proceedings of the 1st Annual PKI Research Workshop*, pages 89–96, April 2002.
- [23] R. Sandhu, V. Bhamidipati, and Q. Munawer. The ARBAC97 model for role-based administration of roles. *ACM Transactions on Information and Systems Security*, 2(1):105–135, February 1999.
- [24] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, February 1996.
- [25] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, November 1979.
- [26] V. Shoup. Practical threshold signatures. In *Advance in Cryptology - EUROCRYPT 2000*, pages 207–220, May 2000.
- [27] D. R. Stinson. *Cryptography: Theory and Practice*. CRC, Boca Raton, 1st edition, 1995.
- [28] X. Wang. Intrusion-tolerant password-enabled PKI. In *Proceedings of the 2nd Annual PKI Research Workshop*, pages 44–53, Gaithersburg, MD, USA, April 28–29 2003. Natl Inst. of Standards and Technology.

- [29] X. Wang, Y. Huang, Y. Desmedt, and D. Rine. Enabling secure on-line DNS dynamic update. In *Proceedings of the 16th Annual Computer Security Applications Conference*, pages 52–58, New Orleans, Louisiana, USA, December 11-15 2000. IEEE Computer Society Press.
- [30] T. Wu. The secure remote password protocol. In *Proceedings of the 1998 Network and Distributed System Security Symposium*, pages 97–111, 1998.
- [31] T. Wu, M. Malkin, and D. Boneh. Building intrusion tolerant applications. In *Proceedings of the 8th USENIX Security Symposium*, pages 79–91, 1999.

Role Sharing in Password-Enabled PKI

Xunhua Wang, Samuel Redwine

James Madison University

{wangxx, redwinst}@jmu.edu

Content

- Background
 - Role and role sharing
 - Password-enabled PKI
- Related work
- Role sharing in password-enabled PKI
 - Threshold role sharing
 - Role sharing for general access structure
- Conclusion

Presentation Roadmap

- Background
 - Role and role sharing
 - Password-enabled PKI
- Related work
- Role sharing in password-enabled PKI
 - Threshold role sharing
 - Role sharing for general access structure
- Conclusion

Role Sharing In General

- Why role?
 - A semantic unit to describe authority and responsibility
 - Users change more frequently than their role
 - Security policies are based on roles, not users
- Why role sharing?
 - Separation of duty
- Role sharing in conventional PKI
 - Threshold cryptography: expensive and complex
 - Voting-based: need a trust machine

Password-Enabled PKI (1/3)

- Where to store your private keys?
 - On a centralized server, protected by passwords
- An end user holds a password, p , only
 - Nothing else: good usability
 - Support user roaming very well
- Building blocks
 - PAKE: client (p), server (password verification data [PVD]) \Rightarrow secure channel
 - Password-based encryption: $E_p(s)$

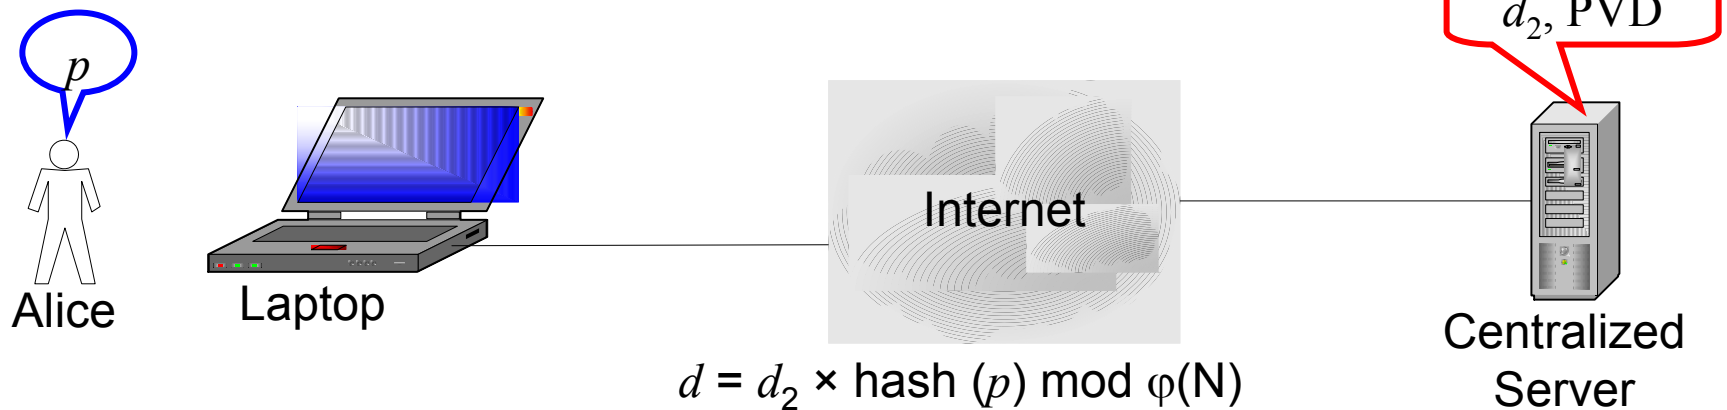

Password-Enabled PKI (2/3)

- Two types of password-enabled PKI
 - Virtual soft token
 - Virtual smartcard
- Use RSA as an example
 - Private key: d
 - Public key: (N, e) , $N = P \times Q$, $\varphi(N) = (P-1) \times (Q-1)$

Password-Enabled PKI (3/3)



Password-Enabled PKI (3/3)



Presentation Roadmap

- Background
 - Role and role sharing
 - Password-enabled PKI
- Related work
- Role sharing in password-enabled PKI
 - Threshold role sharing
 - Role sharing for general access structure
- Conclusion

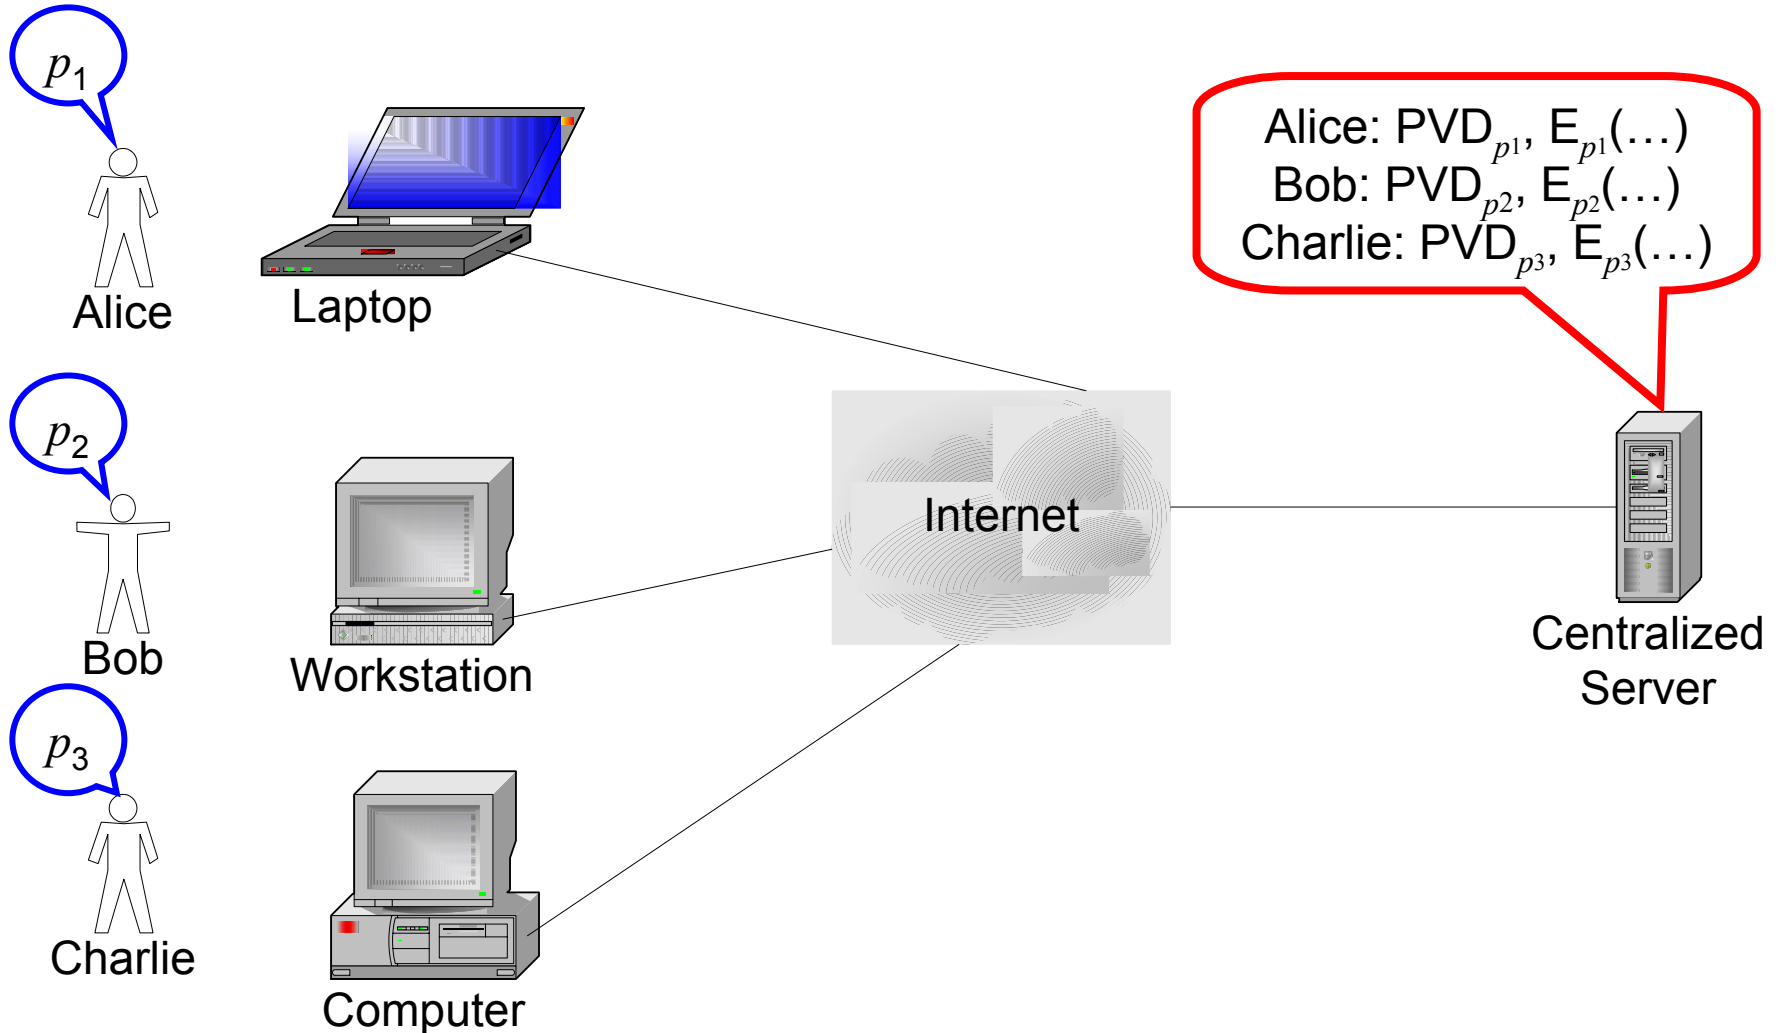
Related Work

- Password-enabled PKI
 - Virtual soft token: (Perlman-Kaufman-99, Kwon-02)
 - Password changes are easy; no server monitoring
 - Virtual smartcard: (Sandhu-Bellare-Ganesan-02)
 - Password changes are hard; allow server monitoring
- Role sharing:
 - Threshold cryptography: not password-enabled
- Server-side PKI activity monitoring: (Boneh-Ding-Tsudik-Wong-01)
 - No password support, not role-oriented
- Our work: password-enabled role sharing
 - Using threshold cryptography primitives and secret sharing schemes

Presentation Roadmap

- Background
 - Role and role sharing
 - Password-enabled PKI
- Related work
- Role sharing in password-enabled PKI
 - Threshold role sharing
 - Role sharing for general access structure
- Conclusion

Password-Enabled Role Sharing



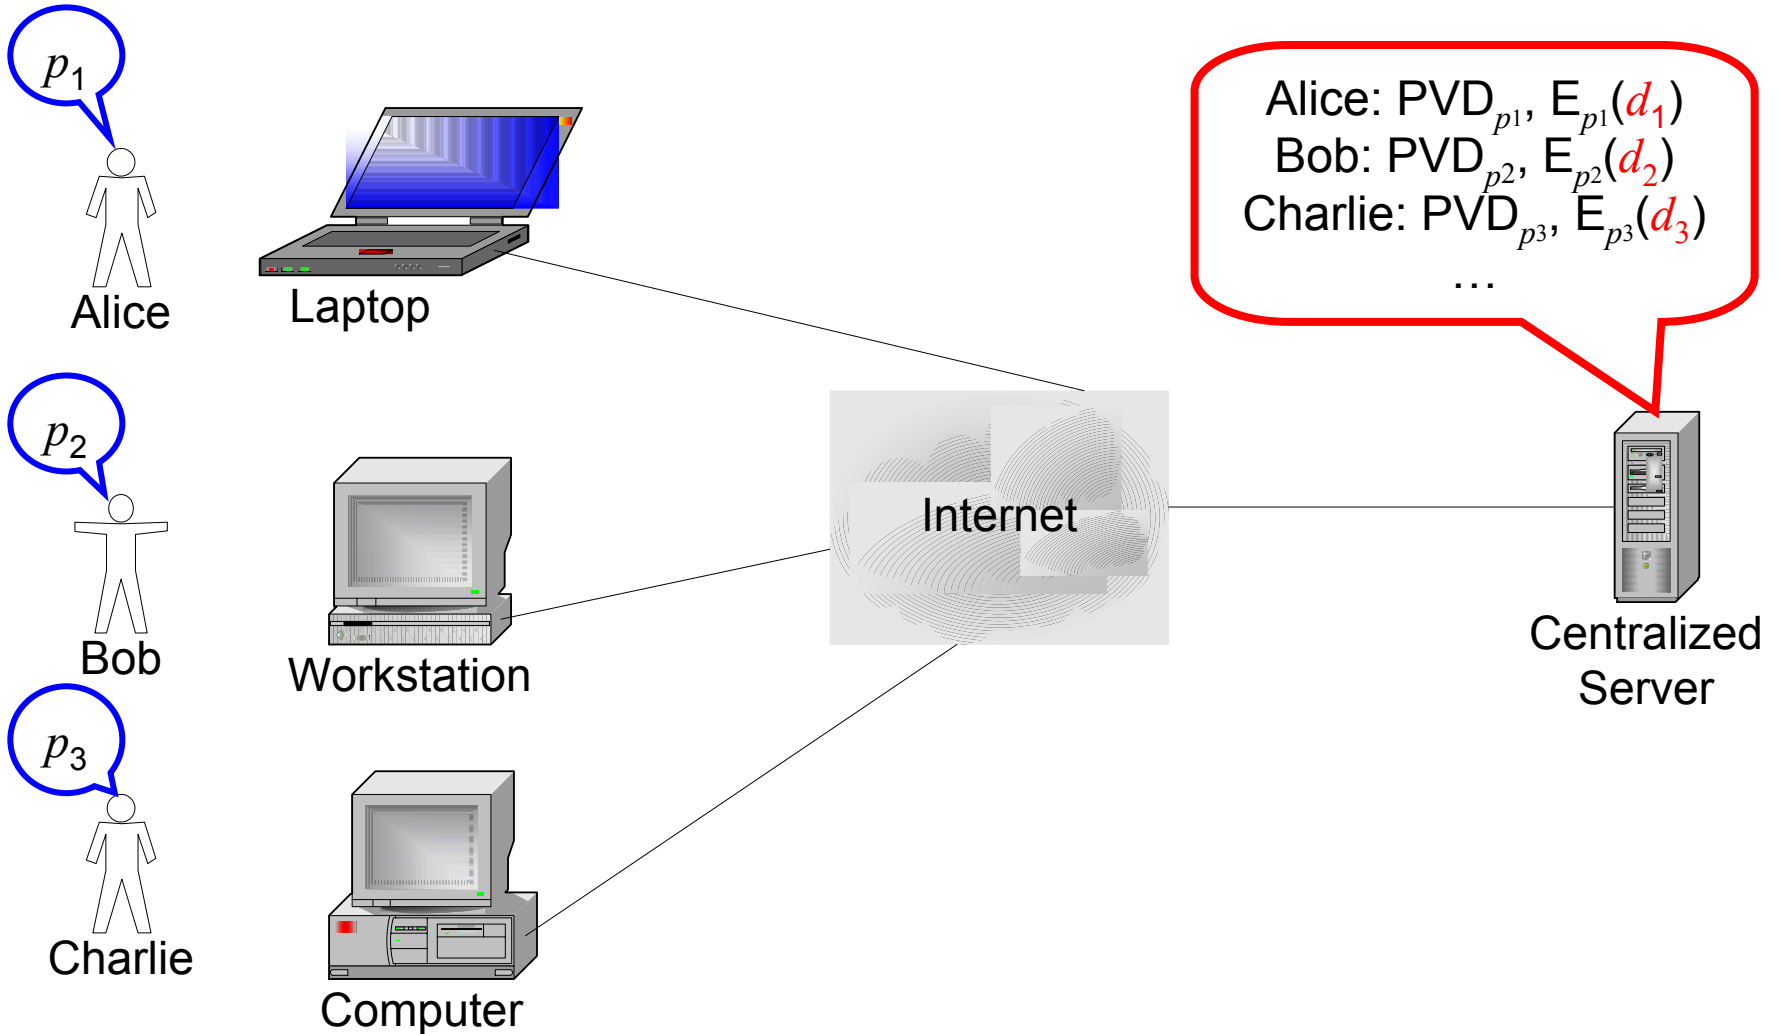
Our Approach

Single-user oriented	Role-oriented
Virtual soft token: (Perlman-Kaufman-99, Kwon-02: easy password change; no server-side monitoring)	Threshold virtual soft token [‡]
Virtual smartcard (Sandhu-Bellare-Ganesan-02: hard password change; allow server-side monitoring)	General access-oriented [‡]
Hybrid password-enabled PKI [‡] (easy password change; allow server-side monitoring)	Threshold [‡]
	General access-oriented [‡]

Threshold Virtual Soft Token

- Role public and private key
 - Public key: (N, e) , $N = p \times q$, $\varphi(N) = (p-1) \times (q-1)$
 - Private key: d
- (t, n) : n users; any t or more can authorize a role-related operation
- Component generation:
 - $d \leftrightarrow_{(t, n)} (d_1, d_2, \dots, d_n) \bmod \varphi(N)$
 - User i is assigned d_i ; $E_{pi}(d_i)$ are stored on the server


Threshold Virtual Soft Token



Virtual Soft Token for General Access Structure

- Not all access structures are threshold-based
- Four users: U_1, U_2, U_3, U_4
 - Authorized subsets $\{U_1, U_2\}, \{U_1, U_3, U_4\}$ can
 - Others, including $\{U_2, U_3, U_4\}$, cannot
- Component generation
 - Benaloh-Leichter-88 (BL88) secret sharing
 - $d_1 + d_2 = d \bmod \varphi(N); d_1 + d_3 + d_4 = d \bmod \varphi(N)$
 - $U_1: d_1; U_2: d_2; U_3: d_3; U_4: d_4$
 - $E_{p_i}(d_i)$ are stored on the server

Hybrid Password-Enabled PKI



- Easy password change: no intensive computations
- Support PKI activity monitoring on the server

Threshold Hybrid Password-enabled PKI

- Component generation:

- $d = d_1 + d_2 \bmod \varphi(N)$

- d_2 is assigned to the server

- $d_1 \leftrightarrow_{(t, n)} (d_{11}, d_{12}, \dots, d_{1n}) \bmod \varphi(N)$

- User i is assigned d_{1i} ; $E_{pi}(d_{1i})$ are stored on the server

Hybrid Password-enabled PKI for General Access Structure

- Four users: U_1, U_2, U_3, U_4
 - Authorized subsets $\{U_1, U_2\}, \{U_1, U_3, U_4\}$ can
 - Others, including $\{U_2, U_3, U_4\}$ cannot
- Component generation:
 - $d = d_1 + d_2 \bmod \varphi(N)$
 - d_2 is assigned to the server
 - $d_{11} + d_{12} = d_1 \bmod \varphi(N); d_{11} + d_{13} + d_{14} = d_1 \bmod \varphi(N)$ {BL88}
 - $U_1: d_{11}; U_2: d_{12}; U_3: d_{13}; U_4: d_{14}$

Operational and Performance Issues

- Password change
 - No intensive computation
- Recovery from password loss
 - Users of an authorized subset can help a user create a new share
- Add and remove a user
- Server compromised?
 - Multiple off-line dictionary attacks
 - One good password is enough to protect the role private key

Conclusion

- Role sharing in password-enabled PKI
- Virtual soft token
 - Threshold access structure
 - General access structure
- Hybrid password-enabled PKI
 - Threshold access structure
 - General access structure
- Good usability

Thank you!

Questions?

Greenpass: Decentralized, PKI-based Authorization for Wireless LANs*

Nicholas C. Goffee, Sung Hoon Kim, Sean Smith, Punch Taylor,
Meiyuan Zhao, John Marchesini

Department of Computer Science/Dartmouth PKI Lab[†]
Dartmouth College
Hanover, New Hampshire USA

Abstract

In Dartmouth's "Greenpass" project, we're building an experimental system to explore two levels of authorization issues in the emerging information infrastructure. On a practical level, we want to enable only authorized users to access an internal wireless network—while also permitting appropriate users to delegate internal access to external guests, and doing this all with standard client software. On a deeper level, PKI needs to be part of this emerging information infrastructure—since sharing secrets is not workable. However, the traditional approach to PKI—with a centralized hierarchy based on global names and heavy-weight X.509 certificates—has often proved cumbersome. On this level, we want to explore alternative PKI structures that might overcome these barriers.

By using SPKI/SDSI delegation on top of X.509 certificates within EAP-TLS authentication, we provide a flexible, decentralized solution to guest access that reflects real-world authorization flow, without requiring guests to download nonstandard client software. Within the "living laboratory" of Dartmouth's wireless network, this project lets us solve real problem with wireless networking, while also experimenting with trust flows and testing the limits of current tools.

1 Introduction

Dartmouth College is currently developing *Greenpass*, a software-based solution to wireless network security in large institutions. Greenpass extends current wireless security frameworks to allow guest access to an institution's wireless network and selected internal resources (as well as to the guest's home system).

This project, which enhances EAP-TLS authentication with SPKI/SDSI-based authorization decisions, is a novel, extensible, feasible solution to an important problem.

- Our solution is **seamless**. Guests can potentially access the same access points and resources that local users can. The same authorization mechanism can apply to

local users, and can also be used for application-level and wired resources.

- Our solution is also **decentralized**: it can accommodate the way that authorization really flows in large academic organizations, allowing designated individuals to delegate network access to guests.

Although we are initially targeting universities, Greenpass may apply equally well in large enterprises.

This paper. This paper provides a snapshot of the current state of our project. Section 2 reviews the problem we seek to solve, and Section 3 reviews the existing wireless security standards we build upon. Section 4 presents weaknesses in some current attempts to secure wireless networks. Section 5 presents our approach: Section 6 discusses the delegation piece; Section 7 discusses the access decision piece. Section 9 discusses future directions, and Section 10 offers some concluding remarks.

More lengthy discussions (e.g., [Gof04, Kim04]) of this work will appear this Spring.

*A previous version of this paper appeared in Dartmouth Dept. of Computer Science technical report TR2004-484

[†]This research has been supported in part by Cisco Corporation, the Mellon Foundation, NSF (CCR-0209144), AT&T/Internet2 and the Office for Domestic Preparedness, Department of Homeland Security (2000-DT-CX-K001). This paper does not necessarily reflect the views of the sponsors. The authors can be reached via addresses of the form `firstname.lastname@dartmouth.edu`

2 The Problem

Wireless network access is ubiquitous at Dartmouth, and we see a future where a lack of wireless network access at a university—including access for visitors to the campus—is as unthinkable as a lack of electricity. Many institutions, however, want to restrict access to their wireless networks for several reasons: the cost of providing network bandwidth and resources; the credibility or liability hit the institution may incur should an outside adversary use the network to launch an attack or spam; the ability to block users who have not installed critical security patches; and the ability (for reasons of license as well as levels-of-protection) to restrict certain local network resources to local users.

Access to a wired network often depends, implicitly, on the physical boundaries of the network. Most establishments do not install courtesy network jacks on the outside walls of their buildings: a standard door, therefore, fulfills most access-control needs. Wireless network traffic, on the other hand, travels on radio waves, extending the network’s physical boundaries. Access control and encryption must be designed into the link layer or higher to prevent unauthorized use and/or eavesdropping.

This future raises some challenges:

- We need to permit authorized local users to access the network.
- We also need to permit selected guests to access the network.
- We must minimize the hassle needed to grant access to guests, and we must accommodate the decentralized ways that authority really flows in large organizations.
- The security should cause little or no additional effort when regular users and guests use the network.
- The type of guests and the manner in which they are authorized will vary widely among the units within an institution.
- We must accommodate multiple client platforms.
- The solution must scale to large settings, more general access policies, and decentralized servers.
- The solution should also extend to *all* authorization—wired or wireless, network or application, guest or intra-institution.
- The solution must be robust against a wide range of failures and attacks.

We have encountered several definitions of “guest access” to a wireless network, many of which differ substantially from our own. Two basic definitions we have seen are as follows:

- **Definition 1.** The trivial solution: the network is open and all passersby, even uninvited ones, can potentially become “guests.”
- **Definition 2.** Insiders can connect to a *VPN (virtual private network)* or to the inside of a firewall, allowing them to access private resources. Guests have basic wireless access—perhaps with a bridge to the Internet—but remain outside the firewall or VPN.

We, on the other hand, want guests to access the inside; that’s the whole point. But we need to control who becomes a “guest.”

We also want to permit authorization to flow the way it flows in the real world; we don’t want a centralized authority (or a central box and rights system purchased from a single vendor) controlling all end-user decisions.

3 Background

Wireless networking comes in two basic flavors. In the *ad hoc* approach, the wireless stations (user devices) talk to each other; in the *infrastructure* approach, wireless stations connect to access points, which usually serve as bridges to a wired network. We are primarily interested in infrastructure networking. Understanding the numerous protocols for access control in infrastructure mode requires wading through an alphabet soup of interrelated standards and drafts. We will provide a brief overview of these standards in this section; Edney and Arbaugh’s recent book [EA03] explores them thoroughly.

Rudimentary access control to a WLAN could be implemented by requiring users to enter the correct *SSID* for the access point they are trying to connect to, or by accepting only clients whose MAC addresses appear on an *access-control list (ACL)*. Both these techniques are easily defeatable, as we discuss below in Section 4.

Wired equivalent privacy (WEP) is a link-layer encryption method offered by the original IEEE 802.11 wireless Ethernet standard. WEP is based on a shared secret between the mobile device and access point. WEP has numerous flaws, which Cam-Winget et al [CWHWW03] and Borisov et al [BGW01] discuss in detail.

WiFi Protected Access (WPA) is a stronger authentication and encryption standard released by the WiFi Alliance, a

consortium of vendors of 802.11-based products. WPA is, in turn, a subset of *802.11i*, the IEEE draft that standardizes future 802.11 security. WPA provides an acceptable security standard for the present, until 802.11i is finalized and becomes widely supported.

WPA and 802.11i both use *802.1x* [CS01], a general access-control mechanism for any Ethernet-based network. 802.1x generalizes the *Extensible Authentication Protocol (EAP)*, originally designed for authentication of PPP dialup sessions.

In a wireless context, 802.1x access control works as follows:

- By trying to connect to an access point, a mobile device assumes the role of *supplicant*.
- The access point establishes a connection to an *authentication server*.
- The access point relays messages back and forth between the supplicant and authentication server. These relayed messages conform to the EAP packet format.
- EAP can encapsulate any of a variety of inner authentication handshakes including challenge-response schemes, password-based schemes (e.g., Cisco's LEAP), Kerberos, and PKI-based methods. The supplicant and authentication server carry out one of these handshakes.
- The authentication server decides whether the supplicant should be allowed to connect, and notifies the access point using an *EAP-Success* or *EAP-Failure* message.

EAP-TLS. Authenticating a large user space suggests the use of public-key cryptography, since that avoids the security problems of shared secrets and the scalability problems of ACLs. One public-key authentication technique permitted within EAP is *EAP-TLS* [AS99, BV98].

TLS (transport layer security) is the standardized version of *SSL (secure sockets layer)*, the primary means for authentication and session security on the Web. In the Web setting, the client and server want to protect their session and possibly authenticate each other. Typically, SSL/TLS allows a Web server to present an X.509 public key certificate to the client and prove knowledge of the corresponding private key. A growing number of institutions (including Dartmouth) also exploit the ability of SSL/TLS to authenticate the client: here, the client presents an X.509 certificate to the server and proves ownership of the corresponding private key. The server can use this certificate to decide whether to grant access, and what Web content to offer. An SSL/TLS

handshake also permits the client and server to negotiate a cryptographic suite and establish shared secrets for session encryption and authentication.

The EAP-TLS variant within 802.1x moves this protocol into the wireless setting. Instead of a Web client, we have the supplicant; instead of the Web server, we have the access point, working in conjunction with the authentication server.

Our approach. WPA with EAP-TLS permits us to work within the existing WiFi standards, but lets the supplicant and access point evaluate each other based on public key certificates and keypairs. Rather than inventing new protocols or cryptography, we plan to use this angle—the expressive power of PKI—to solve the guest authorization problem.

4 Black Hat

As part of this project, we began exploring just how easy it is to examine wireless traffic with commodity hardware and easily-available hacker tools. Right now:

- We can watch colleagues surf the Web and read their email.
- We can read the “secret” (non-broadcast) SSID for local networks.
- We can read the MAC addresses of supplicants permitted to access the network.
- We can tell our machine (Windows or Linux) to use a MAC address of our own choosing (such as one that we just sniffed).

The lessons here include:

- We can easily demonstrate that security solutions which depend on secret SSIDs or authenticated MAC addresses do not work.
- The current Dartmouth wireless network is far more exposed than nearly all our users realize; the paradigm shift from the wired net has substantially changed the security and privacy picture, but social understanding (and policy) lags behind. We suspect this is true of most wireless deployments.

We conjecture that any solution that does not use cryptography derived from entity-specific keys will be susceptible to sniffing attacks and session hijacking.

5 The Overall Approach

We have already built a basic prototype of Greenpass, and are planning a pilot in the near future. Our prototype consists of two basic tools. The first automates the process of issuing credentials to a guest by allowing certain local users to issue *SPKI/SDSI authorization certificates* [EFL⁺99a, EFL⁺99b] to guests. The second tool is a *RADIUS (Remote Authentication Dial In User Service) server* [Rig00, RWC00, RWRS00] that carries out a standard EAP-TLS handshake for authentication, but has been modified to consult SPKI/SDSI certificates for authorization of non-local users. Neither tool requires users to have software beyond what is typically installed on a Windows laptop (covering most of our user space); other platforms need 802.1x supplicant software, which is provided with recent versions of Mac OS X and is readily available for Linux.

Authorization in real life. In the physical world, a guest gets access to a physical resource because, according to the local policy governing that resource, someone who has the power to do so said it was OK. In a simple scenario, Alice is allowed to enter lab because she works for Dartmouth; guest Gary is allowed to come in because Alice said it was OK. Gary's authorization is decentralized (Dartmouth's President Wright doesn't know about it) and temporary (it vanishes tomorrow). More complex scenarios also arise in the wild: e.g., Gary may only have access to certain rooms, and it must be Alice (and not Bob, since he doesn't work on that project) who says OK.

For a wireless network in a large institution, the decision to grant authorization will not always be made by the same Alice—and may in fact need to reflect policies and decisions by many parties. PKI can handle this, by enabling verifiable chains of assertions.

Authorization in EAP-TLS. EAP-TLS specifies a way for a RADIUS server to *authenticate* a client, but leaves open the specification of *authorization*. Often, a RADIUS server will allow any supplicant to connect who authenticates successfully—i.e., whose certificate was signed by a CA the RADIUS server has been configured to trust. This approach does not adequately reflect real-life authorization flow as just described. Alice can see to it that Gary, her guest, obtains access to the wireless network, but she must do so by asking a central administrator to issue Gary an X.509 certificate from Dartmouth's own CA. It is possible for the RADIUS server to trust multiple CAs, such as those of certain other universities, but this option remains inflexible if a guest arrives from an institution not recognized by the existing configuration. (Another option that merits further investigation, however, is the possibility of

linking RADIUS servers using more advanced trust path construction and bridging techniques. One such implementation is the Trans-European Research and Education Networking Association's (TERENA) [TER] top-level European RADIUS server, a hierarchy of RADIUS servers connecting the Netherlands, the UK, Portugal, Finland, Germany, and Croatia.)

Conceivably, a RADIUS server could perform any of a number of authorization checks between the time that a supplicant authenticates successfully and the time that the RADIUS server transmits an EAP success or failure code. In other words, we can modify a RADIUS server to base its decision on some advanced authorization scheme. Policies could be defined by policy languages such as XACML; or by signed *authorization certificates* as defined by *Keynote* [BFIK99] and its predecessor *PolicyMaker* [BFL96], by the *X.509 attribute certificate (AC)* standard [FH02], and by SPKI/SDSI.

SPKI/SDSI. For our Greenpass prototype, we settled on SPKI/SDSI for three main reasons: (1) it focuses specifically on the problem we are trying to solve (authorization), (2) its central paradigm of *delegation* gives us precisely the decentralized approach to guest access we desire, and (3) its lightweight syntax makes it easy to process and to code for.

SPKI/SDSI differs from a traditional X.509-based PKI in three important ways:

- SPKI/SDSI uses public keys as unique identifiers: people and other *principals* are referred to as holders of particular public keys, rather than as entities with particular names.
- A SPKI/SDSI certificate binds an authorization directly to a public key, rather than binding a name to a public key (as an X.509 certificate does) or an authorization to a name (as an ACL entry or attribute certificate typically does).
- Any person or entity, not just a dedicated CA, can potentially issue a SPKI/SDSI certificate. In particular, the recipient of a SPKI/SDSI authorization can optionally be authorized to *delegate* his privilege to further users.

SPKI/SDSI therefore solves some of the problems with guest authorization. First, even if a guest's home organization issued him an X.509 certificate, we cannot use it to authenticate the guest (i.e., bind the guest to a unique identifier) if the issuer of the certificate is not trusted. A SPKI/SDSI authorization certificate, however, binds an authorization to a particular keyholder without an intermediate

naming step: therefore, we can bind credentials to a guest's public key. The public key acts as the sole identifying information for the guest ("authentication," then, means proving ownership of the key).

Additionally, SPKI/SDSI delegation provides a straightforward way to implement guest access. Dartmouth can issue Alice, a professor, a SPKI/SDSI certificate granting her the ability to delegate access to guests.¹ If Alice invites Gary to campus to deliver a guest lecture, he will probably request access to the network. Alice can simply issue Gary a short-lived SPKI/SDSI certificate (vouching for the public key in his X.509 certificate) that grants him access to the network while he is on campus. No central administrator need be contacted to fulfill Gary's request.

Alternative approaches to authorization. Other approaches to delegated guest access are available, and each has its own balance of advantages and disadvantages.

An X.509 AC can grant a short-lived authorization to the holder of a particular X.509 identity certificate, in much the same way as we use SPKI/SDSI. Attribute certificates address the problem of authorization, but are intended to be issued by small numbers of *attribute authorities* (AAs); the attribute certificate profile [FH02] states that chains of ACs are complex to process and administer, and therefore recommends against using them for delegation. Doing so would amount to emulating SPKI/SDSI's functionality using attribute certificates. If standard 802.11 clients were able to transmit attribute certificates along with identity certificates as part of the EAP-TLS handshake, ACs would have deserved more attention, as they would have provided a convenient means to transmit authorization information to a RADIUS server.

The PERMIS system [COB03, Per] allows end users to issue authorization certificates, but this feature is intended to support a multiple, fixed set of authorization certificate issuers rather than a delegation model.

We also could have implemented guest access by placing temporary ACL entries in a central database. "Delegation" could be implemented by allowing authorized delegators to modify certain portions of the ACL. Ultimately, however, would like to support a "push" model of delegation where guests carry any necessary credentials and present them upon demand, allowing us to further decentralize future versions of Greenpass (see Section 9). Decentralizing authorization policies using signed certificates also eliminates the

¹In our current scheme, Alice uses an X.509 certificate, signed by the local CA, to gain access to the network herself; she must obtain a SPKI/SDSI certificate only if she needs to delegate to a guest without a locally-signed X.509 certificate.

need for a closely-guarded machine on which a central ACL is stored.

We chose SPKI/SDSI because it reflects, in our minds, the most straightforward model of real-world delegation. A comparison of alternative approaches would provide a worthwhile direction for future work.

6 Delegation

Assume that a new guest arrives and already holds an X.509 identity certificate containing a public key. In order to obtain wireless connectivity, the guest must obtain a SPKI certificate that conveys the privilege of wireless network access directly to his own public key.

To obtain this certificate, the guest will find a local user who can delegate to him (e.g., the person who invited him in the first place). This *delegator* must then learn the guest's public key. This step requires an information path from the guest's machine to the delegator's. Once the delegator learns the guest's key, he can issue a SPKI certificate with the guest as its subject.

We also need a way to ensure that the key the delegator authorizes to use the wireless network is really the key held by the guest. Otherwise, an adversary might inject his own public key into the communication channel between guest and delegator, tricking the delegator into authorizing the adversary instead of the intended guest. Dohrmann and Ellison describe a nearly identical problem in *introducing* the members of a collaborative group to one another [DE02]. Their solution was to display a *visual hash* of the public key being transferred on both the keyholder's device and the recipient's device: this allows the recipient to quickly compare the two visual images, which should appear identical if and only if the recipient received a public key value identical to the one stored on the originating device. We adopted this same approach; further details are given below.

Guest interface. We chose to use a Web interface to allow a guest to introduce his public key to a delegator. Web browsers are ubiquitous: we can safely assume that any user who wishes to access our network will have a Web browser installed. This technique gives us an advantage over, e.g., infrared transfer, wired transfer, or passing of some storage medium, all of which might be incompatible with certain client devices.

Both our delegation tool and our modified RADIUS server rely on the observation that standard X.509 certificates, and standard SSL/TLS handshakes (including EAP-TLS) perform three functions that we need:

- An X.509 certificate contains the value of its owner's public key.
- An SSL/TLS handshake *presents* an X.509 certificate (and thus the owner's public key value).
- An SSL/TLS handshake, if it succeeds, also *proves* that the authenticating party owns the private key corresponding to the subject public key in the presented X.509 certificate.

With this observation, it becomes clear that, if the guest's Web browser supports SSL/TLS client authentication, then he can present his public key value to a Web site using this functionality.

When a guest arrives, therefore, he must connect to our Web application to present his existing X.509 certificate. Therefore, he must obtain *some* wireless connectivity even before he is authorized. We are experimenting with various ways to enable this by creating an "unauthorized" VLAN (for newly-arrived guests) and an "authorized" VLAN (for local users and authorized guests); we present our approach in more detail in Section 7.

When a guest connects to our Web application, he will see a welcome page helping him through the process of presenting his certificate. We handle three situations at this point:

- If the guest's Web browser presents an SSL client certificate, we allow the option of presenting it immediately.
- We also allow the guest to upload his certificate from a PEM-formatted file on his local disk. (Browser and OS keystore tools usually allow a user to export his X.509 certificate as a PEM file. Since the purpose is to transfer the certificate to another user, a PEM file typically does *not* contain the user's private key.)
- If the guest does not already have a keypair and certificate, he can connect to a "dummy" CA page (separate from the main Dartmouth CA) that lets him generate a keypair and obtain a temporary X.509 certificate. (This should not be a standard approach, because a proliferation of client keypairs impairs usability. Note that the sole purpose of the dummy CA is to get the guest a keypair—we are therefore exempt from standard CA worries such as securing the registration process and protecting the CA's private keys.)

We implemented the guest interface using simple CGI scripts served by an Apache web server. Our installation of Apache includes *mod_ssl*, which we configure to request (but not require) SSL client authentication. (We had to set a

seldom-used option in *mod_ssl* that forces Apache to accept the guest's certificate even if it was signed by an unknown CA. Our purpose here is to learn a stranger's public key, not to authenticate a known user.) Therefore, if the guest has installed a client certificate in his Web browser, it will present it to our Web server. Our CGI scripts use OpenSSL to process the guest's X.509 certificates.

The dummy CA uses the standard enrollment functionality included in Web browsers that support SSL client authentication. The guest visits the CA page and enters (possible fake) identifying information. The page includes code that, when he submits the identifying information, causes his Web browser to generate a keypair, store the private key, and submit the public key to our Web server. The dummy CA then issues a new X.509 certificate back to the guest's Web browser, which stores it in its keystore. We support both the IE and the Netscape/Mozilla methods of enrollment.

After the guest presents his X.509 certificate by one of the above methods, our Web server generates a visual hash of it using the *Visprint* [Gol, Joh] program. (This program transforms the MD5 hash of an object into an image using IFS fractals.)

After the guest uploads his certificate using one of the above methods, our Web server stores it in a temporary repository from which the delegator can retrieve it.

Delegator interface. A delegator first visits the same Web server as the guest, and searches for the guest's X.509 certificate by entering pieces of identifying information such as the guest's name and organization. After this step, the delegator verify the certificate's authenticity and construct and sign a SPKI certificate.

Signing a SPKI certificate is problematic, because it requires access to the delegator's private key. A private key must be well-protected so that adversaries cannot use it to sign data that did not actually originate from the owner. Software usually signs data of a very specific type (email, Word documents, authentication challenges, certificates) to prevent misuse of the key.

We therefore needed to build a special software tool for signing SPKI certificates. We considered a number of alternative ways to implement this, including a custom application which delegators would have to download, but for the prototype, we settled on using a trusted Java applet (screenshot shown in Figure 1). Trusted applets are hashed and signed by an entity that the user of the applet trusts, ensuring that the applet has not been modified to do anything the signing entity did not intend. Sun's Java plugin for Web browsers, by default, gives trusted applets greater privileges




Figure 1: A screenshot of our delegator tool (a trusted Java applet, shown running under Mac OS 10.3). Before delegation, the delegator has to verify the identity of the guest’s public key using a visual hash comparison. Also notice the inputs for validity interval and whether or not to allow further delegation.

than standard applets, including the ability to access the local filesystem on the client machine. Our applet can, therefore, load the delegator’s private key from a local file² and, after prompting for a password to decrypt the key, use it to sign a SPKI certificate.

Our Web server generates a page with a reference to the delegation applet, and provides the guest’s PEM-encoded certificate as an argument to the applet. The applet uses standard Java cryptography functionality to extract the public key from this certificate, and uses a Java SPKI/SDSI library from MIT [Mor98] to construct and sign a SPKI certificate that delegates wireless access privileges to the guest. The applet allows the delegator to specify a validity interval for the new certificate and choose whether or not the recipient should be able to delegate further. We have almost finished porting the Visprint code to Java so we can build the visual hash verification step into the applet as well.

7 Making the Decision

We now consider the process by which our modified RADIUS decides whether to admit users.

7.1 The decision process

Local users. In the initial case, local users show authorization (via EAP-TLS) by proving knowledge of a private

²The applet prompts the delegator to choose an appropriate keystore file the first time it is run, and saves its location to a local preferences file for future signing sessions. We currently support PKCS12 keystore files. In the future, we would like to support various platforms’ OS keystores.

key matching an X.509 identity certificate issued by the local CA. Once the TLS handshake succeeds, the supplicant is granted access. On most platforms, the supplicant must choose which certificate to submit only on the first successful attempt; the machine will remember which certificate to use on subsequent attempts, making the authentication process transparent to local users.

Guests. Authorized guests also authenticate via EAP-TLS using an X.509 certificate. (In this case, “authentication” consists only of proving knowledge of the private key, since we cannot trust the certificate’s naming information.) The RADIUS server uses a different process, however, to decide whether the user is authorized. It must find a valid SPKI/SDSI certificate chain originating from a principal it trusts that ultimately grants access privileges to the supplicant’s public key.

In preliminary sketches, we also involved the delegator’s X.509 certificate, but that does not seem to be necessary. As a consequence, the delegator doesn’t necessarily need to have a centrally-issued X.509 identity certificate; we consider this further in Section 9.

The algorithm. Putting it all together, the modified RADIUS server follows the following procedure, illustrated by the flowchart in Figure 2:

- The supplicant initiates an EAP-TLS authentication handshake.
- If the supplicant cannot present an identity certificate, we shunt them to a special VLAN on which the supplicant can only connect to our delegation tool’s “welcome” page.
- If the supplicant *can* present an identity certificate, we then evaluate it as follows:
 - If the certificate is valid and issued by the local CA, then we accept it.
 - Otherwise, if we can obtain and verify a valid SPKI/SDSI chain supporting it, we accept it.
 - Otherwise, we reject the certificate and shunt the supplicant to our “welcome” page.
- If we accept the certificate, and the supplicant proceeds to prove knowledge of the private key, then we let him in.
- Otherwise, we shunt the supplicant to our “welcome” page.

This procedure modifies standard EAP-TLS implementations only by changing how the server decides to accept a given supplicant certificate.




Figure 2: Decision flowchart used by the RADIUS server. If the supplicant is a local Dartmouth user (i.e., presents an X.509 certificate issued by the Dartmouth CA), then the supplicant only needs to prove knowledge of the private key associated with the certificate. Otherwise, if the supplicant is a guest, the RADIUS server checks for a SPKI certificate chain vouching for the supplicant’s public key.

Getting the certificates (“pull” approach). To carry out the guest user case, the RADIUS server needs to know the X.509 identity certificate, the public key of whatever source-of-authority SPKI/SDSI chains will originate from, and the SPKI/SDSI certificate chain itself. EAP-TLS gives us the first, and we can build in the second. But how do we find the relevant SPKI/SDSI certificates?

One solution would be to have the delegation process leave the authorization certificates in a reliable, available directory where servers can access them; since the data is self-validating, maintenance of this directory should be automatic. When the RADIUS server needs to verify a guest’s SPKI/SDSI credentials, it can “pull” up the credentials it requires from the directory. We can organize these certificates as a forest: guest authorization certificates are children of the delegation certificates that signed them.

- The source-of-authority tool needs to write new dele-

gator certificates to this directory.

- The delegator tool needs to read delegator certificates from this directory, and write new guest authorization certificates back.
- The RADIUS server needs to be able to ask for delegator-authorization chains whose leaves speak about a given public key.

The directory itself can perform time-driven checks for expiration.

Our implementation currently uses the “pull” approach just described: SPKI/SDSI certificates are maintained in a cache that the RADIUS server can query via XML-RPC. The RADIUS server queries the cache about a particular public key; the cache itself finds a chain, if it exists, verifies it, and returns it. (To make our prototype more secure, we

need to use authenticated XML-RPC messages or move the decision procedure onto the same machine as the RADIUS server.)

Getting the certificates (“push” approach). The centralized solution above is somewhat unsatisfying, because it introduces a centralized component (even if this component does not have significant security requirements). It would be slicker to find a way for the delegator and guest themselves to carry around the necessary certificates, since the necessary information paths will exist. When necessary, the guest can “push” the necessary credentials to the RADIUS server for validation.

We note that HTTP cookies will provide most of the functionality we need. (We will add a message to the guest welcome page notifying users of what browser features will need to be enabled, including cookies and Java, in order to use our services.)

- The delegator will be interacting with the source-of-authority signing tool when their delegation certificate is created; the delegation certificate could be saved at the delegator machine as a cookie.
- At delegation time, both the delegator and the guest will be interacting with the delegation tool. The tool can read the delegator’s certificate as a cookie, and then store that and the new authorization certificate as cookies as the guest’s machine.

The only remaining question would be how to get these two cookies from the guest machine to the RADIUS server, when an authorized guest connects. One approach would be to add a short-term SPKI/SDSI store to the RADIUS server. When deciding whether to accept an X.509 certificate not issued by the Dartmouth CA, the server looks in this store for a SPKI/SDSI certificate chain for this X.509 cert. If none can be found, the supplicant is routed to a special Web page, that will pick up their two certificate cookies (this requires the guest must have a browser running) and save them in the store.

In this decentralized approach, it also might make sense to have the delegation tool save newly created SPKI/SDSI chains in the short-term store at the RADIUS server, since the guest will likely want to use the network immediately after being delegated to.

Changing VLANs. We now have two scenarios—when first receiving delegation, and in the above decentralized store approach—where a supplicant will be connected through the access point to the special VLAN, but will want

to then get re-connected to the standard network. In both scenarios, the guest will be interacting with the Web server we have set up on the special VLAN.

One way to handle this would be for our server to display a page telling the guest how to cause their machine to drop the network and re-associate. However, this is not satisfying, from a usability perspective.

Instead, it would be nice to have our server (and back-end system) cause this action automatically. One approach would be to use the administrative interface provided by the access point. For example, the Cisco 350 access point (that we’re experimenting with) permits an administrator, by a password-authenticated Web connection, to dis-associate a specific supplicant (after which the supplicant re-initializes the network connection, and tries EAP-TLS again). We could write a daemon to perform this task, when it receives an authenticated request from our backend server. The server needs to know *which* access point the supplicant is associated with; however, in both scenarios, the RADIUS server has recently seen the supplicant MAC and access point IP address, since it told the access point to route this supplicant down the special VLAN. If nothing else, we can cache this information in a short-term store that the daemon can query.

We plan to explore other approaches here as well.

7.2 Executing the decision

On the server side, we are currently using FreeRadius version 0.9.2, running on a Dell P4 workstation running Red Hat 9, and an Apache web server running on another Dell P4 workstation running Red Hat 9. We’re testing with a Cisco 350 access point, with a Cisco Catalyst 2900 series XL switch and a hub to connect the two machines running the RADIUS server and Web server.

Setup. In our prototype, we have the access point configured to provide two different SSIDs. The broadcast SSID is called “Guest user” and authentication is not needed. It associates all users onto VLAN 2, the guest VLAN. The SSID “Dartmouth user” is not broadcast, and requires EAP authentication. Supplicants who pass EAP authentication are associated to this SSID on VLAN 1, the native VLAN that has access to the whole network. (We will abbreviate these designations as V_1 and V_2 in the following discussion.)

Our VLAN configuration is illustrated in Figure 3. The RADIUS server is connected to V_1 on the switch and the Web server is connected to V_2 . The access point, connected

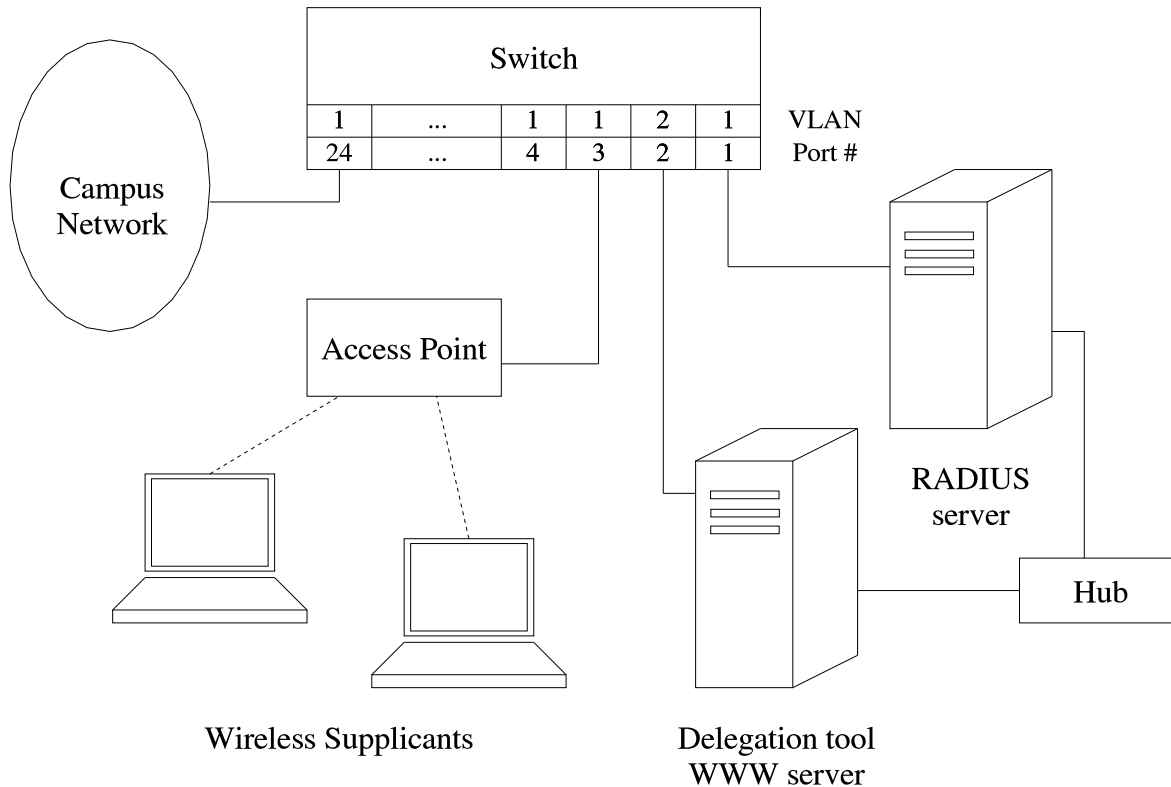


Figure 3: The setup of the Greenpass prototype. The switch is configured to associate VLANs with physical port numbers. The Web server is the only element currently housed on VLAN 2. Eventually, VLAN trunking will be used to communicate between the RADIUS server and the web server, eliminating the need for the private connection that exists between the two.

to V_1 , is configured to query the RADIUS server for user authentication. The hub connects the two machines and allows them to communicate to one another through the resulting private connection. In the future, we will purchase a router capable of VLAN trunking, which will allow the Web server to exist on both VLANs; this will eliminate the need for the private connection through a hub.

Configuration. The EAP-TLS module of FreeRADIUS uses OpenSSL to execute the SSL/TLS handshake between the supplicant and the RADIUS server. After changing the appropriate RADIUS configuration files to enable EAP-TLS authentication and linking the OpenSSL libraries [Sul02], the RADIUS server was ready to accept EAP-TLS authentication attempts. The client file was configured to only accept requests sent from an access point (called a *network authentication server*, NAS, in the RADIUS protocol) with a Dartmouth IP address and the user file was set to only allow EAP (in our case EAP-TLS) authentication for all users, placing the user on V_1 if successfully authenticated. A shared secret between the RADIUS server and the NAS secures communication between these two components.

In order to use EAP-TLS authentication, the RADIUS server

needs a trusted root CA so that it knows which certificates to accept. The RADIUS server also needs its own server certificate and key pair issued by the trusted root CA for authenticating itself to the supplicant in the handshake process. Local users are given a key pair and issued client certificates signed by the trusted root CA. OpenSSL can be used to generate key pairs, create a root certificate, and issue server and client certificates [Ros03]. Once the RADIUS server has a trusted root CA to refer to, it can handle authentication requests from the access point.

We modified the RADIUS server code to link with XML-RPC libraries we installed on the same machine. These libraries allow the RADIUS server to communicate with the cache, mentioned above, that stores SPKI/SDSI certificates and searches for chains authorizing a given principal to connect.

The decision process. The RADIUS server idles and waits for packets. When it receives an EAP Access-Request packet, it checks to see if the NAS that sent the packet is recognized and the shared secret is correct. If so, then it looks at the packet and sees what type of authentication is used.

Since the SSID is configured to require EAP authentication, the RADIUS server should only receive EAP authentication requests from the NAS.

Once the EAP-TLS module is done executing, the decision to accept or reject the supplicant has already been made and is packed into the response packet. Thus it is necessary to intercept the EAP-TLS module before a reject decision is made to accommodate any modifications to the decision process.

Our modification determines if there is an error code returned by reading the supplicant's certificate. For example, the most common case would be the certificate is issued by an unrecognized CA. Once the validity checks are finished, we read the resulting error code to see if the validation passed or failed. If it passed, then the certificate presumably was issued by the known CA and the supplicant has provided knowledge of the corresponding private key. If the handshake failed due to an unrecognized CA, however, we use XML-RPC to query the Java SDSI library code about the public key of the X.509 certificate provided. The library uses the SPKI/SDSI certificate chain discovery algorithm proposed by Clark et al [CEE⁺01]. If the Java code finds a valid SPKI certificate chain vouching for the supplicant, then we accept the supplicant and the EAP-TLS module returns an accept code. If such a SPKI/SDSI certificate chain cannot be found, then the user is rejected. Once graceful VLAN switching is implemented, the unauthorized guest will be placed on V_2 and see a web browser window with instructions for obtaining guest access.

8 Related Work

Balfanz et al [BDS⁺03] propose using secret keys to let wireless parties authenticate. We've already noted related work [DE02] in the "introduction" problem between two devices.

In the SPKI/SDSI space, the Geronimo project at MIT [Cla01, May00] uses SPKI/SDSI to control objects on an Apache Web server. The project uses a custom authorization protocol, with an Apache module handling the server side of the protocol and a Netscape plug-in handling the client side. The protocol can be tunneled inside an SSL channel for further protection; the authors also considered replacing X.509 with SPKI/SDSI within SSL. Koponen et al [KNRP00] propose having an Internet cafe operator interact with a customer via infrared, and then having that customer authenticate to the local access point via a SPKI/SDSI certificate; however, this work does not use standard tools and institution-scale authentication servers.

Canovas and Gomez [CG02] describe a distributed management system for SPKI/SDSI name certificates and authorization certificates. The system contains name authorities (NAs) and authorization authorities (AAs) from which entities can request name and authorization certificates, including certificates which permit the entity to make requests of further NAs and RAs. The system takes advantage of both name certificates that define groups (i.e., roles) and authorization certificates that grant permissions to either groups or individual entities.

9 Future Directions

Initially, we plan to "take the duct tape" off of our current prototype, and try it in a more extensive pilot. Beyond this initial work, we also hope to expand in several directions.

No PKI. We note that our approach could also accommodate the scenario where *all* users are "guests" with no keypairs—in theory, obviating the need for an X.509 identity PKI for the local population. For example, if an institution already has a way of authenticating users, then they could use a modified delegator tool that:

- authenticates the delegator (via the legacy method)
- sees that the delegator has a self-signed certificate (like our guest tool does)
- then signs a SPKI/SDSI delegator certificate for this public key (like our delegator tool does).

In some sense, the division between the X.509 PKI and the delegated users is arbitrary. It would be interesting to explore the implications of dividing the population in other ways than users versus guests (perhaps "permanent Dartmouth staff" versus "students likely to lose their expensive smart-card dongles while skiing").

Not just the network. Many types of digital services use X.509 identity certificates as the basis for authentication and authorization. For example, at Dartmouth, we're migrating many current Web-based information services to use X.509 and client-side SSL/TLS. In the Greenpass pilot, we're adding flexibility to wireless access by extending X.509/TLS with SPKI/SDSI. This same PKI approach can work for networked applications that expect X.509, such as our Web-based services.

In the second phase, we will extend the Greenpass infrastructure to construct a single tool that allows delegation of authorization to networked applications as well as to the network itself.

Not just EAP-TLS. Some colleagues insist that *virtual private networks* with client-side keypairs are the proper way to secure wireless networks. In theory, our scheme should work just as well there. In the second phase, we plan to try this.

Alternative approaches to hash verification. An attacker could potentially abuse our delegator applet if the delegator chooses to skip the fingerprint-verification step. Visual fingerprints are designed to discourage users from skipping crucial verification steps: it is faster and less painful to compare two visual fingerprints than to compare hashes represented as hexadecimal strings. We must devise either a method which ensures that the delegator cannot skip this step,³ or a method that takes humans out of the loop entirely. Balfanz et al. [BSSW02] suggest an introduction phase based on a *location-limited channel*; this approach might allow us to eliminate human interaction from the introduction phase in the future. We are also considering alternative models of fingerprint verification: for example, using PGPfone's [PGP] mapping of hash values to word lists would allow introduction to take place over the phone as well as in person.

Location-aware authorization and services. By definition, the RADIUS server making the access-control decision knows the supplicant's current access point. In some scenarios, we may want users to access the network only from certain access points; in some scenarios, users should be able to access some applications only from certain access points; potentially, the nature of the application content itself may change depending on access location.

In the second phase, we plan to extend the Greenpass infrastructure to enable authorization certs to specify the set of allowable access points. We will also enable the RADIUS back-end to sign short-term certificates testifying to the location of the supplicant (which requires an authorization cert for the server public key), and to enable applications to use these certificates for their own location-aware behavior. For example, we might put different classes of users (professors, students, guests, etc.) on different VLANs according to the resources we would like them to access. It might also be interesting to allow certain users to access the WLAN only from certain locations—e.g., conference rooms and lecture halls.

Who is being authorized? Campus environments are not monolithic. At Dartmouth, we already have multiple schools, departments, and categories of users within departments. Managing authorization of such internal users is a

³An in-progress revision of our delegator tool requires the user to select the correct visual hash from among several choices.

vexing problem. Centralized approaches are awkward and inflexible: a colleague at one university ended up developing over 100 different user profiles; a colleague at another noted she has to share her password to team-teach a security course, because the IT department has no other way to let her share access to the course materials.

In the second phase, we plan to extend the Greenpass infrastructure to support authorization delegation for "local users" as well as guests, and to permit local users to easily manage authorization for information resources they own or create.

Devices. Currently, laptops are probably the most common platform for access to wireless networks. Other platforms are emerging, however. At Dartmouth, students and staff already carry around an RFID tag embedded in their ID cards, a research team is developing experimental wireless PDAs for student use, and we are beta-testing Cisco's new VoIP handset device; we're also testing Vocera's device for WiFi voice communication.

In the second phase, we plan to explore using these alternate devices in conjunction with Greenpass. For example, a department's administrative assistant might be able to create a SPKI/SDSI cert and enter it in a directory simply by pointing a "delegation stick" (RFID tag reader) at the student (detecting the student's ID card). In another example, when a physician at the Dartmouth-Hitchcock Medical Center collars a passing colleague for advice on a difficult case, he might be to delegate permission to read that file simply by pointing his PDA at the colleague's PDA.

Distributed authorization. The PKI community has long debated the reason for PKI's failure to achieve its full potential. The technology exists and has clear benefits; adoption and deployment has been a challenge.

One compelling hypothesis is that the centralized, organizational-specific hierarchy inherent in traditional approaches to PKI, compounded by a dependence on usable, globally unique names and awkward certificate structure, did not match the way that authorization really flows in human activities. By permitting authorization to start at the end-users (rather than requiring that it start at centralized places), and by using a system (SPKI/SDSI) designed to address the namespace and structure issues, Greenpass may overcome these obstacles.

In the second phase, we plan to extend Greenpass to reproduce real-world policies more complex than just "Prof. Kotz said it was OK," to examine (with our colleagues in the Dept of Sociology) how readily this authorization system matches

the norms of human activity, and to examine whether humans are able to manage the user interfaces our Greenpass tools provide.

We also plan to take a closer look at how other authorization schemes might fit in this setting, in comparison to SPKI/SDSI. Some candidates include the X.509-based PERMIS attribute certificate system might work in this setting [COB03, Per], as well as KeyNote [BFIK99, Key]. Nazareth [Naz03] gives an overview of many such systems.

10 Conclusion

In this paper we described a method of securing a wireless network while providing meaningful guest access. We added a step to EAP-TLS authentication that performs an additional authorization check based on SPKI/SDSI certificates. By using SPKI/SDSI, we eliminate the need for a cumbersome central authority; by grafting it on top of the existing X.509-based PKI, we do not require our users to install any additional client software.

The two major components of the Greenpass project are the delegation tools and the modified RADIUS server. The delegation tools automate the process of creating temporary SPKI/SDSI certificates for a guest, allowing an authorized (but not necessarily computer-savvy) delegator to grant an invited guest permission to use the network. The modified RADIUS server takes into account that guests will want to access the network and checks for guest credentials before making a decision to accept or reject a supplicant's request for network access.

The goal of our project is to create a solution that implements delegation in a way that reflects real-world authorization flow that does not rely too heavily on a centralized authority; SPKI/SDSI allows us to accomplish this goal. Our future work will allow us to investigate how our solution fits with other existing ideas, hopefully resulting in a solution that is secure, completely decentralized, and capable of adapting to new technology and delegation policies.

Acknowledgments

We gratefully acknowledge Kwang-Hyun Baek, Bob Brentrup, Bill Cote, Peter Glenshaw, Dave Kotz, Brad Noblet, and our colleagues at Cisco for the advice, and Eileen Ye for her initial SPKI/SDSI explorations. We plan to make code for this project available under an open-source license.

References

- [AS99] Bernard Aboba and Dan Simon. PPP EAP TLS Authentication Protocol. IETF RFC 2716, October 1999.
- [BDS⁺03] Dirk Balfanz, Glenn Durfee, Narendar Shankar, Diana Smetters, Jessica Staddon, and Hao-Chi Wong. Secret Handshakes from Pairing-Based Key Agreements. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pages 180–196, May 2003.
- [BFIK99] Matt Blaze, Joan Feigenbaum, John Ioannidis, and Angelos D. Keromytis. The KeyNote Trust-Management System Version 2. IETF RFC 2704, September 1999.
- [BFL96] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized Trust Management. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 1996.
- [BGW01] Nikita Borisov, Ian Goldberg, and David Wagner. Intercepting Mobile Communications: The Insecurity of 802.11. In *Proceedings of Mobicom 2001*, 2001.
- [BSSW02] Dirk Balfanz, D. K. Smetters, Paul Stewart, and H. Chi Wong. Talking to Strangers: Authentication in Ad-Hoc Wireless Networks. In *Proceedings of the Network and Distributed System Security Symposium*, February 2002.
- [BV98] Larry J. Blunk and John R. Vollbrecht. PPP Extensible Authentication Protocol (EAP). IETF RFC 2284, March 1998.
- [CEE⁺01] D. Clark, J. Elien, C. Ellison, M. Fredette, A. Morcos, and R. Rivest. Certificate Chain Discovery in SPKI/SDSI. *Journal of Computer Security*, 9(4):285–322, 2001.
- [CG02] Oscar Canovas and Antonio F. Gomez. A distributed credential management system for spki-based delegation scenarios. In *Proceedings of the 1st Annual PKI Research Workshop*, April 2002.
- [Cla01] Dwaine E. Clarke. SPKI/SDSI HTTP Server / Certificate Chain Discovery in SPKI/SDSI. Master’s thesis, Massachusetts Institute of Technology, September 2001.
- [COB03] David W. Chadwick, Alexander Otenko, and Edward Ball. Role-Based Access Control with X.509 Attribute Certificates. *IEEE Internet Computing*, March-April 2003.
- [CS01] IEEE Computer Society. IEEE Standard for Local and metropolitan area networks: Port-Based Network access control. IEEE Standard 802.1X-2001, October 2001.
- [CWHWW03] Nancy Cam-Winget, Russ Housley, David Wagner, and Jesse Walker. Security Flaws in 802.11 Data Link Protocols. *Communications of the ACM*, 46(5):35–39, May 2003.
- [DE02] Steve Dohrmann and Carl M. Ellison. Public-Key Support for Collaborative Groups. In *Proceedings of the 1st Annual PKI Research Workshop*, April 2002.
- [EA03] Jon Edney and William A. Arbaugh. *Real 802.11 Security*. Addison-Wesley, 2003.
- [EFL⁺99a] Carl M. Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian M. Thomas, and Tatu Ylonen. Simple public key certificate. IETF Internet Draft, draft-ietf-spki-cert-structure-06.txt, July 1999.
- [EFL⁺99b] Carl M. Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian M. Thomas, and Tatu Ylonen. SPKI Certificate Theory. IETF RFC 2693, September 1999.
- [FH02] Stephen Farrell and Russell Housley. An Internet Attribute Certificate Profile for Authorization. IETF RFC 3281, April 2002.
- [Gof04] Nicholas C. Goffee. Greenpass Client Tools for Delegated Authorization in Wireless Networks. Master’s thesis, Dartmouth College, May 2004. (expected).
- [Gol] Ian Goldberg. Visual Fingerprints (Visprint software homepage). <http://www.cs.berkeley.edu/~iang/visprint.html>.
- [Joh] David Johnston. Visprint, the Visual Fingerprint Generator. <http://www.pinkandaint.com/oldhome/comp/visprint/>.
- [Key] Keynote home page. <http://www.cis.upenn.edu/~keynote/>.

- [Kim04] Sung Hoon Kim. Greenpass RADIUS Tools for Delegated Authorization in Wireless Networks. Master's thesis, Dartmouth College, May 2004. (expected).
- [KNRP00] Juha Koponen, Pekka Nikander, Juhana Rasanen, and Juha Paajarvi. Internet Access through WLAN with XML-encoded SPKI Certificates. In *Proceedings of NORDSEC 2000*, 2000.
- [May00] Andrew J. Maywah. An Implementation of a Secure Web Client Using SPKI/SDSI Certificates. Master's thesis, Massachusetts Institute of Technology, May 2000.
- [Mor98] Alexander Morcos. A Java Implementation of Simple Distributed Security Architecture. Master's thesis, Massachusetts Institute of Technology, May 1998.
- [Naz03] Sidharth Nazareth. SPADE: SPKI/SDSI for Attribute Release Policies in a Distributed Environment. Master's thesis, Department of Computer Science, Dartmouth College, May 2003.
<http://www.cs.dartmouth.edu/~pkilab/theses/sidharth.pdf>.
- [Per] Permis home page.
- [PGP] PGPfone: Pretty Good Privacy Phone Owner's Manual. <http://web.mit.edu/network/pgpfone/manual/>.
- [Rig00] Carl Rigney. RADIUS Accounting. RFC 2866, June 2000.
- [Ros03] Ken Roser. HOWTO: EAP-TLS Setup for FreeRADIUS and Windows XP Supplicant. <http://3w.denobula.com:50000/EAPTLS.pdf>, February 2003.
- [RWC00] Carl Rigney, Ward Willats, and Pat R. Calhoun. RADIUS Extensions. RFC 2869, June 2000.
- [RWRS00] Carl Rigney, Steve Willens, Allan C. Rubens, and William Allen Simpson. Remote Authentication Dial In User Service (RADIUS). RFC 2865, June 2000.
- [Sul02] Adam Sulmicki. HOWTO on EAP/TLS authentication between FreeRADIUS and XSupplicant. <http://www.missl.cs.umd.edu/wireless/eaptls/>, April 2002.
- [TER] Trans-european research and education networking association.
<http://www.terena.nl/tech/task-forces/tf-mobility/>.

Greenpass: Decentralized, PKI-Based Authorization for Wireless LANs

*Nicholas C. Goffee, Sung Hoon Kim, Sean Smith,
Punch Taylor, Meiyuan Zhao, John Marchesini*

**Department of Computer Science
Dartmouth College**

April 12, 2004



Motivation

Our wireless security goals

- Allow *internal* access by *authorized guests*
- Without requiring custom client software

Our PKI goals

- Move away from *centralized, name-based hierarchies*
- Accommodate *real-world trust flow*

This talk

Background

- The WPA wireless authentication standard
- Current approaches to authorization and guest access

Decentralized guest authorization

- Authorization certificates
- Modified authentication/authorization server

Delegation process

- Guest introduction
- Delegator tool
- Decentralization

Next steps

WPA authorization and guest access

WPA itself doesn't specify an authorization procedure

Centralized schemes *without* guest access

- Allow all users certified by the “right” CA
- Authenticate, then consult ACL

Centralized schemes *with* guest access

- Trust multiple CAs (with or without ACL)
- Allow unauthenticated users/guests ***outside*** the firewall

Decentralized authorization

A SPKI/SDSI authorization certificate

- Binds *authorization* → *public key**
- Can optionally allow *delegation*
- Delegation allows decentralized guest authorization

Example SPKI/SDSI certificate (unsigned)

```
(cert
  (issuer (hash md5 |BuFWyi13EpqzJtMff8DcsA==|))
  (subject (hash md5 |9WgBTLBGk6kIIvJVwZLbAg==|))
  (propagate)
  (tag (greenpass-pilot-auth))
  (valid (not-after "2004-07-02_17:43:06")))
```

WLAN authorization with SPKI/SDSI (ideal)

- User presents credential (cert) to AP
- Unfortunately, standard client software doesn't support this

Authentication/authorization procedure

RADIUS server

- Modified to look up SPKI/SDSI chain for guests

Local users

- X.509 certificates signed by local CA
- Standard EAP-TLS handshake will succeed

Guests

- RADIUS server uses EAP-TLS to extract public key
- Doesn't recognize issuer CA in guest's X.509 certificate
- Checks **certificate store** for relevant SPKI/SDSI chain
- Authorized guests get access
- Unauthorized guests are put on **restricted VLAN**


Delegation process

Guest introduces public key to delegator

- Guest connects to Web server on restricted VLAN
- Web server gets (or generates) guest's X.509 certificate
- Web server displays guest's *visual fingerprint*

Delegator signs new authorization certificate

- Delegator connects to same Web server
- Uses visual fingerprint to verify guest's identity
- Delegator generates and signs new SPKI/SDSI certificate
- Delegator signature generated by *trusted Java applet*
- Sends fresh certificate to certificate store



Greenpass guest introduction

Thank you for introducing yourself. Your **subject ID number** is **3707**; please give this number to your delegator. Your delegator will also ask to see the **visual fingerprint** displayed below in order to verify your identity.



After your delegator has signed an **authorization certificate** granting you access to our network, you should click on the **Continue** button to install it in your web browser.*

*Your authorization certificate will be stored as a cookie within your Web browser.



Greenpass delegation tool

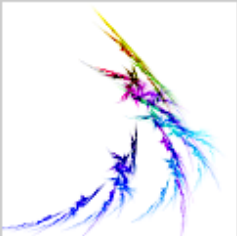

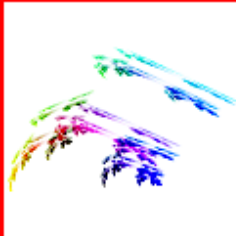
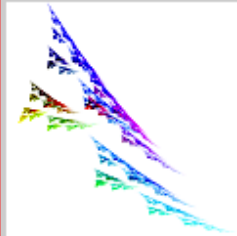
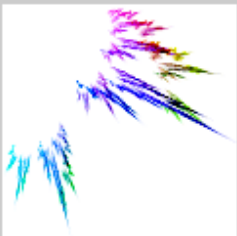
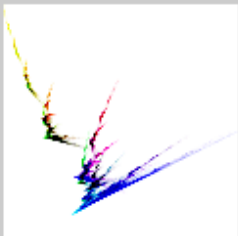
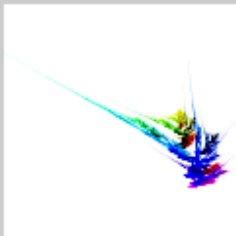
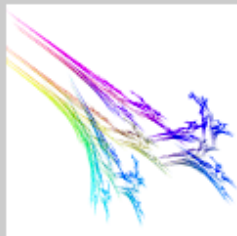
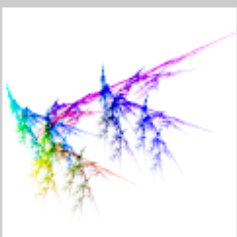
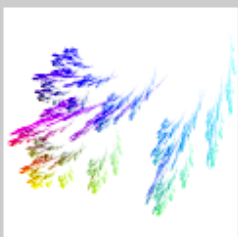
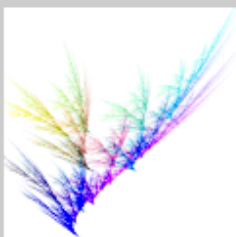
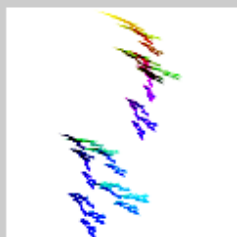
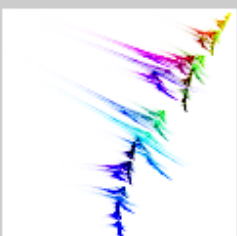
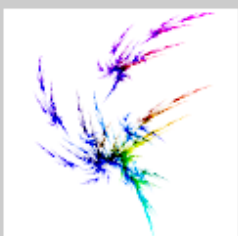
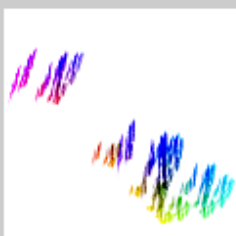
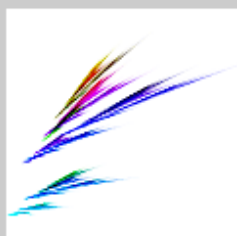
Java Applet Window

Delegation options

Valid for 2 Days

Propagate

To verify that you delegate to the correct person, please choose the visual fingerprint below that matches the fingerprint shown on your guest's screen.


			
			
			
			

Help... Cancel Delegate!

Decentralization

Certificate store has become a short-term certificate cache

- After delegation, guest can pick up certificate chain
- Chain gets stored as a cookie on guest's machine
- Guest can use chain to reauthorize without introduction process



Greenpass front page

Your status is: REAUTHORIZED_USER

Next steps

Take off the duct tape

- Decentralize certs by using HTTP cookies (**done**)
- Move to router that can handle VLAN trunking
- Test various OS/wireless card configurations
- Move from FreeRADIUS to Cisco ACS

Try it in the real world

- Is SPKI/SDSI sufficiently expressive?
- What about SAML or PERMIS or proxy certificates or...?
- What about revocation?*

Next steps (cont'd)

Apply our tools to other settings

- VPN
- Application-level resources
- Location-sensitive policy

Try other PKI models

- Move the hybrid X.509/SPKI boundary
- No X.509 at all
- Cross-domain

Acknowledgments

Funding:

- Cisco Corporation
- Mellon Foundation
- National Science Foundation
- AT&T/Internet2
- Department of Homeland Security

Collaborators:

- Dr. Sean Smith (advisor)
- Sung Hoon Kim (RADIUS code)
- Kimmy Powell (testing, pilot)
- Meiyuan Zhao (SPKI/SDSI libraries)
- Kwang-Hyun Baek (wireless security details)
- Punch Taylor (network configuration)
- John Marchesini (black hat)

X.509 Proxy Certificates for Dynamic Delegation

Von Welch¹, Ian Foster^{2,3}, Carl Kesselman⁴, Olle Mulmo⁵, Laura Pearlman⁴, Steven Tuecke², Jarek Gawor², Sam Meder³, Frank Siebenlist²

¹National Center for Supercomputing Applications, University of Illinois

²Mathematics and Computer Science Division, Argonne National Laboratory

³Department of Computer Science, University of Chicago

⁴Information Sciences Institute, University of Southern California

⁵Royal Institute of Technology, Sweden

Contact author: Von Welch (vwelch@ncsa.uiuc.edu)

Abstract

Proxy credentials are commonly used in security systems when one entity wishes to grant to another entity some set of its privileges. We have defined and standardized X.509 Proxy Certificates for the purpose of providing restricted proxying and delegation within a PKI-based authentication system. We present here our motivations for this work coming from our efforts in Grid security, the Proxy Certificate itself, and our experiences in implementation and deployment.

1 Introduction

“Grids” [10] have emerged as a common approach to constructing dynamic, inter-domain, distributed computing and data collaborations. In order to support these environments, Grids require a light-weight method for dynamic delegation between entities across organizational boundaries. Examples of these delegation requirements include granting privileges to unattended processes which must run without user intervention, the short-term sharing of files for collaboration, and the use of brokering services which acquire resources (e.g., storage, computing cycles, bandwidth) on behalf of the user.

The Globus Toolkit[®] [11] has emerged as the dominant middleware for Grid deployments worldwide. The Grid Security Infrastructure (GSI) [39,2,9] is the portion of the Globus Toolkit that provides the fundamental security services needed to support Grids. GSI provides libraries and tools for authentication and message protection that use standard X.509 public key certificates [5,16], public key infrastructure (PKI), the SSL/TLS protocol [6], and X.509 Proxy Certificates, an extension defined for GSI to meet the delegation requirements of Grid communities.

Proxy Certificates allow an entity holding a standard X.509 public key certificate to delegate some or all of its privileges to another entity which may not hold X.509 credentials at the time of delegation. This delegation can be performed dynamically, without the assistance of a third party, and can be limited to arbitrary subsets of the delegating entity’s privileges. Once acquired, a Proxy Certificate is used by its bearer to authenticate and establish secured connections with other parties in the same manner as a normal X.509 end-entity certificate.

Proxy Certificates were first prototyped in early implementations of GSI. Subsequently, they have been refined through standardization in the IETF PKIX working group [17]

and have achieved RFC status. (At the time of this writing, the Proxy Certificate internet draft [37] has passed IETF-wide public comment and is only awaiting assignment of an RFC number). GSI currently implements this standard.

GSI and Proxy Certificates have been used to build numerous middleware libraries and applications that have been widely deployed in large production and experimental Grids [2,3,4,19,35]. This experience has proven the viability of proxy delegation as a basis for authorization within Grids, and has further proven the viability of using X.509 Proxy Certificates.

We start with a discussion of the requirements that spurred our use of X.509 public key certificates and motivated our development of Proxy Certificates. We follow with a technical description of the format of Proxy Certificates in Section 3. Section 4 describes how Proxy Certificates can be used to achieve single sign-on and delegation, and Section 5 describes how Proxy Certificates can be integrated with different types of authorization systems. Section 6 discusses current implementations and applications of Proxy Certificates. Section 7 discusses performance issues with Proxy Certificates and security tradeoffs in regards to those issues. We conclude with a discussion of related work in Section 8 and a summary in Section 9.

2 Motivation

We discuss first our motivation for the use of X.509 certificates and PKI as the basis for our GSI implementation. Then we discuss the motivations that lead to the creation of Proxy Certificates as an enhancement to standard X.509 public key certificates.

2.1 Motivation for X.509 Certificates

GSI uses X.509 public key certificates and Secure Socket Layer (SSL) for authentication not only because these are well-known technologies with readily available, well-tested open source implementations, but because the trust model of X.509 certificates allows an entity to trust another organization's certification authority (CA) without requiring that the rest of its organization do so or requiring reciprocation by the trusted CA.

This flexibility of trust model for X.509 certificates was a deciding factor between X.509 certificates and other common authentication mechanisms. For example, Kerberos [29] requires that all cross-domain trust be established at the domain level, meaning that organizations have to agree to allow cross-domain authentication, which can often be a heavy-weight administrative process. In many common Grid deployments, only a few users and resources at a particular organization may participate in the Grid deployment, making the process of acquiring buy-in from the organization as a whole to establish the authentication fabric prohibitive.

2.2 Motivation for Proxy Certificates

The establishment of X.509 public key certificates and their issuing certification authorities provides a sufficient authentication infrastructure for persistent entities in Grids. However, several use cases exist that are not well covered by X.509 public key certificates alone.

- *Dynamic delegation*: It is often the case that a Grid user needs to delegate some subset of their privileges to another entity on relatively short notice and only for a brief amount of time. For example, a user needing to move a dataset in order to use it in a computation may want to grant to a reliable file transfer service the necessary rights to access the dataset and storage so that it may perform a set of file transfers on the user's behalf. Since these actions may be difficult to predict, having to arrange delegation ahead of time through some administrator is prohibitive.
- *Dynamic entities*: In addition to delegation to persistent services and entities, the requirement exists to support delegation of privileges to services that are created dynamically, often by the user them self, that do not hold any form of identity credential. A common scenario is that a user submits a job to a computational resource and wants to delegate privileges to the job to allow it to access other resources on the user's behalf, for example, to access data belonging to the user on other resources or start sub-jobs on other resources. An important point here is that the user wants to delegate privileges specifically to the job and not to the resource as a whole (i.e., other jobs being run by other users on the resource should not share the rights).
- *Repeated Authentication*: It is common practice to protect the private keys associated with X.509 public key certificates either by encrypting them with a pass phrase (if stored on disk) or by requiring a PIN for access (if on a smart card). This technique poses a burden on users who need to authenticate repeatedly in a short period of time, which occurs frequently in Grid scenarios when a user is coordinating a number of resources.

A number of existing mechanisms could satisfy the first use case. For example, user-issued X.509 attribute certificates [8] could be used to delegate rights to other bearers of X.509 public key certificates. However, the heavy-weight process of vetting associated with the issuing of public key certificates makes it prohibitive to use this method for the dynamically created entities described in the second use case: acquiring public key certificates for dynamically created, and often short-lived entities, would be too slow for practical use. It would have been possible to use other means for authenticating these dynamic entities, for example bare keys as described in Section 8.5, but this approach would have required protocol modifications (or a new protocol) to accommodate the new authentication mechanism.

The third scenario could be solved by caching the pass phrase or PIN required for access to the private key. However, this caching increases the risk of compromising the private key if the memory storing the pass phrase or PIN is somehow accessible to an attacker or is written out to disk (e.g., if it is swapped or in a core dump). In addition, reliably caching the PIN for a set of simultaneously running applications is a non-trivial software engineering exercise.

These requirements led us to develop an authentication solution that allows users to create identities for new entities dynamically in a light-weight manner, to delegate privileges to those entities (again in a dynamic, light-weight manner), to perform single sign-on, and that allows for the reuse of existing protocols and software with minimal

modifications. The result is the X.509 Proxy Certificate, which we describe in the following sections.

We note that while it may be possible to use Proxy Certificates for uses other than authentication, delegation and message protection, for example the signing or encryption of long-lived documents, these alternate uses were not motivating factors in the Proxy Certificate design and we have not investigated such use .

3 Description of Proxy Certificates

We now describe the contents of a Proxy Certificate and briefly discuss methods of revocation and path validation.

3.1 Proxy Certificate Contents

Proxy Certificates use the format prescribed for X.509 public key certificates [5,16] with the prescriptions described in this section on the contents. Proxy Certificates serve to bind a unique public key to a subject name, as a public key certificate does. The use of the same format as X.509 public key certificates allows Proxy Certificates to be used in protocols and libraries in many places as if they were normal X.509 public key certificates which significantly eases implementation.

However, unlike a public key certificate, the issuer (and signer) of a Proxy Certificate is identified by a public key certificate or another Proxy Certificate rather than a certification authority (CA) certificate. This approach allows Proxy Certificates to be created dynamically without requiring the normally heavy-weight vetting process associated with obtaining public key certificates from a CA.

The subject name of a Proxy Certificate is scoped by the subject name of its issuer to achieve uniqueness. This is accomplished by appending a CommonName relative distinguished name component (RDN) to the issuer's subject name. The value of this added CommonName RDN should be at least statistically unique to the scope of the issuer. The value of the serial number in the Proxy Certificate should also be statistically unique to the issuer. Uniqueness for both of these values in our implementations is achieved by using the hash of the public key as the value. Unique subject names and serial numbers allow Proxy Certificates to be used in conjunction with attribute assertion approaches such as attribute certificates [8] and have their own rights independent of their issuer.

The public key in a Proxy Certificate is distinct from the public key of its issuer and may have different properties (e.g., its size may be different). As we describe in more detail in Section 4, except when using Proxy Certificates for single sign-on, the issuer does not generate the public key-pair and has no access to the private key.

All Proxy Certificates must bear a newly-defined critical X.509 extension, the Proxy Certificate Information (PCI) extension. In addition to identifying Proxy Certificates as such, the PCI extension serves to allow the issuer to express their desire to delegate rights to the Proxy Certificate bearer and to limit further Proxy Certificates that can be issued by that Proxy Certificate holder.

The issuer's desires towards delegation to the Proxy Certificate bearer are expressed in the PCI extension using a framework for carrying policy statements that allow for this delegation to be limited (perhaps completely disallowed). There exist today a number of policy languages for expressing delegation policies (e.g., Keynote, XACML, XrML), instead of defining a new mechanism or selecting a single existing policy language for expressing delegation policy (which probably would have bogged the process of standardizing Proxy Certificates down considerably), Proxy Certificates instead allow the issuer to use any delegation policy expression it chooses. The only restriction being that the issuer needs to know (through some out-of-band method) that the relying party understands its method of expression. This allows different deployments to select (or create) a method of delegation policy expression best suited for their purposes.

This use of arbitrary policy expressions is achieved through two fields in the PCI extension: a policy method identifier and a policy field. The policy method identifier is an object identifier (OID) that identifies the delegation policy method used in the policy field. The policy field then contains an expression of the delegation policy that has a format specific to the particular method (and may be empty for methods that do not require additional policy). For example, the identifier could contain an OID identifying the method as XACML and then the policy would contain an XACML policy statement.

The Proxy Certificate RFC defines two policy methods that must be understood by all implementations of Proxy Certificates (in addition to any more sophisticated methods they may implement):

- *Proxying*: This policy type indicates that the issuer of the Proxy Certificate intended to delegate all of their privileges to the Proxy Certificate bearer.
- *Independent*: This policy type indicates that the issuer of the Proxy Certificate intended the Proxy Certificate by itself to convey none of the issuer's privileges to the bearer. In this case the Proxy Certificate only serves to provide the bearer with a unique identifier, which may be used in conjunction with other approaches, such as attribute certificates, to grant its bearer privileges.

For both of these methods, the policy field is empty since the intended delegation policy is explicit in the type.

Certificate attribute	X.509 Public key certificate	X.509 Proxy Certificates
Issuer/Signer	A certification authority	A public key certificate or another Proxy Certificate
Name	Any as allowed by issuer's policy	Scoped to namespace defined by issuer's name
Delegation from Issuer	None	Allows for arbitrary policies expressing issuer's intent to delegate rights to Proxy Certificate bearer.
Key pairs	Uses unique key pair	Uses unique key pair

Table 1: Comparison of X.509 public key certificates and X.509 Proxy Certificates.

The PCI extension also contains a field expressing the maximum path lengths of Proxy Certificates that can be issued by the Proxy Certificate in question. A value of zero for this field prevents the Proxy Certificate from issuing another Proxy Certificate. If this field is not present, then the length of the path of Proxy Certificates, which can be issued by the Proxy Certificate, is unlimited.

3.2 Proxy Certificate Path Validation

Validation of a certificate chain has two distinct phases. First validation of the certificate chain up to the public key certificate occurs, as described by RFC 3280 [16]. Validation of the Proxy Certificate portion of the chain is then performed as described in the Proxy Certificate RFC [37]. In summary these rules are:

- Ensuring each Proxy Certificate has a valid Proxy Certificate Information extension as described in the previous section;
- Each Proxy Certificate must have a subject name derived from the subject name of its issuer;
- Verifying the number of Proxy Certificates in the chain does not exceed the maximum length specified in any of the Proxy Certificate Information extensions in the chain; and
- Storing the delegation policies of each Proxy Certificate so that the relying party can determine the set of rights delegated to the bearer of the end Proxy Certificate used to authenticate.

3.3 Revocation of Proxy Certificates

There currently exists no implemented method for revocation of Proxy Certificates. The intent is that Proxy Certificates are created with short life spans, typically on the order of hours (with eight hours being the default of our implementation). Therefore, revocation has not been a pressing issue since this short lifetime limits the length of misuse if a Proxy Certificate were to be compromised. However, Proxy Certificates can be uniquely identified in the same manner as normal end-entity certificates, through the issuer and serial number, so the potential exists to revoke them using the same mechanisms (e.g., CRLs [16] or OCSP [28]).

4 Use for Single Sign-on and Delegation

In this section we describe how Proxy Certificates can be used to perform single sign-on and delegation.

4.1 Enabling Single Sign-on

Normally the private key associated with a set of long-term X.509 credentials is protected in some manner that requires manual authentication on the behalf of its owner. While this process serves to provide a high level of protection of the private key, it can be prohibitively burdensome if the user needs to access the key frequently for authentication to other parties.

Proxy Certificates solve this problem by enabling single sign-on: that is, allowing the user to manually authenticate once in order to create a Proxy Certificate which can be used repeatedly to authenticate for some period of time without compromising the protection on the user's long-term private key. This is accomplished by creating a new key pair (composed of a public and private key), and by subsequently using the user's long-term private key to create a short-lived Proxy Certificate. The Proxy Certificate binds the new public key to a new name and delegates some or all of the user's privileges to the new name. The Proxy Certificate and the new private key are then used by the bearer to authenticate to other parties. Since the Proxy Certificate has a short lifetime, it is typically permissible to protect it in a less secure manner than the long-term private key. In practice this means the Proxy Certificate private key is stored on a local file system and is protected by only local file system permissions, which allows the user's applications to access it without any manual intervention by the user.

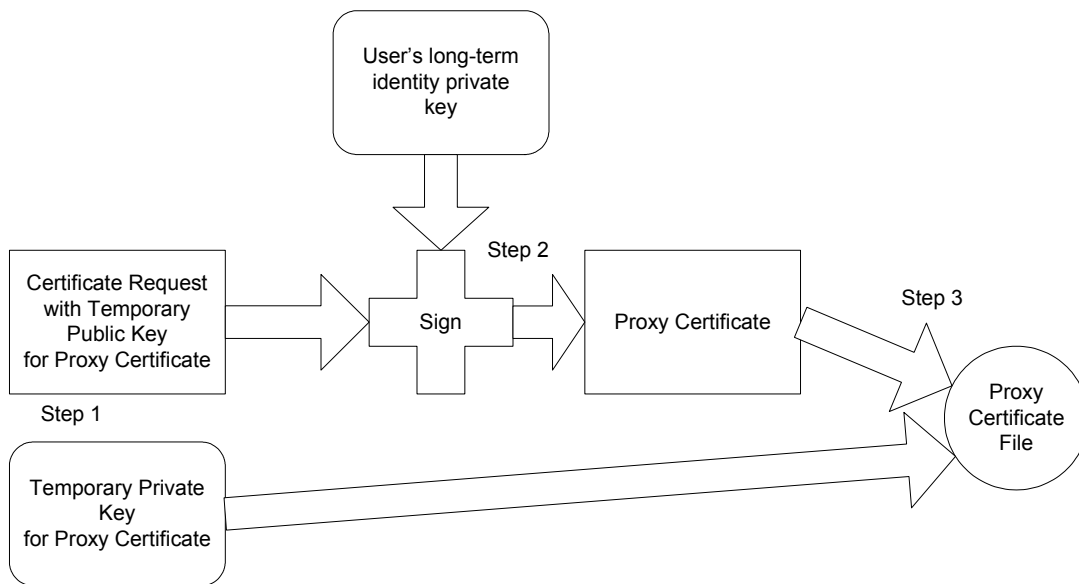


Figure 1: Creation of a Proxy Certificate for single sign-on. Steps are described in the text.

The process of creating a Proxy Certificate for single sign-on is shown in Figure 1. The steps, which are normally all done by a single application run by the user, are:

1. A new key pair, consisting of a public and private key, is generated for use in the Proxy Certificate. The public key is encoded in a certificate request [20] for further processing.
2. The user's private key associated with their long-term public key certificate is accessed (possibly requiring the manual entering of a pass phrase or PIN by the user) to sign the certificate request containing the public key of the newly generated key pair hence generating a Proxy Certificate. After signing the Proxy Certificate, the user's long-term private key can remain secured (or the associated smart card can be removed) until the Proxy Certificate expires.
3. The Proxy Certificate and its associated private key are then placed in a file. This file is protected only by local file system permissions to allow for easy access by the user.

When the Proxy Certificate expires, this process is repeated by the user to generate a new key pair and Proxy Certificate. The result from the perspective of the user is that manual authentication is required only infrequently to enable applications to authenticate on their behalf.

4.2 Delegation over a Network

Proxy Certificates can also be created so as to delegate privileges from an issuer to another party over a network connection without the exchange of private keys. This delegation process requires that the network connection be integrity-protected to prevent malicious parties from tampering with messages, but does not require encryption as no sensitive information is exchanged.

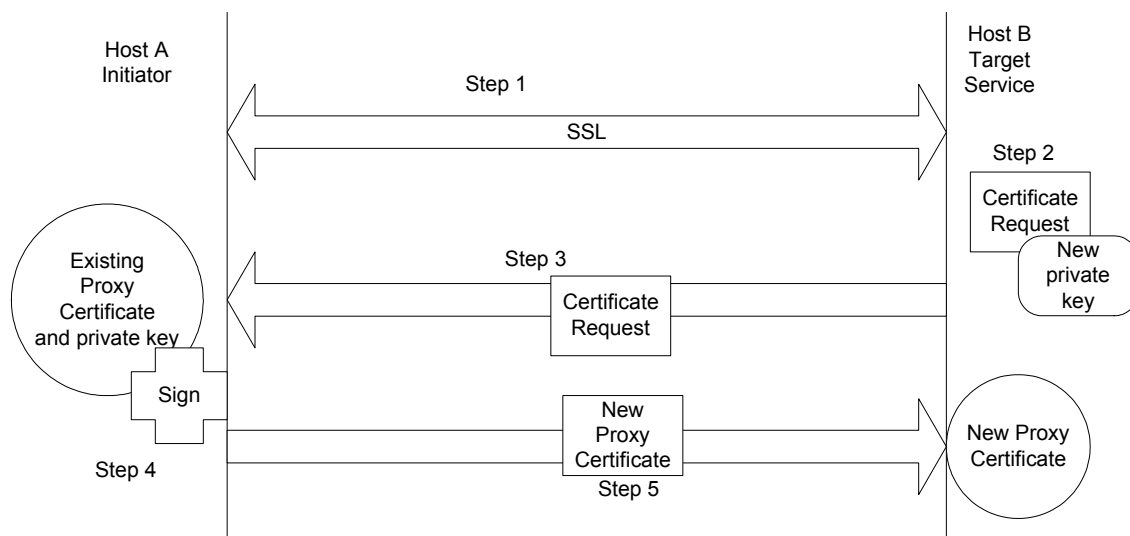


Figure 2: Delegation of a Proxy Certificate over a secured network connection. Steps are described in the text.

Figure 2 shows the steps involved in the delegation of privileges by creation of a Proxy Certificate over a network connection:

1. The initiator, on host A at left, connects to the target service on host B at right. The initiator and target service perform mutual authentication, the initiator using its existing Proxy Certificate and the target service uses the public key certificate of its own (not shown). After authentication, an integrity protected channel is established. These two steps can be accomplished by using the SSL protocol.
2. After the initiator expresses its desire to delegate by some application-specific means, the target service generates a new public and private key pair.
3. With the new public key, a signed certificate request is created and sent back over the secured channel to the initiator.
4. The initiator uses the private key associated with its own Proxy Certificate to sign the certificate request, generating a new Proxy Certificate containing the newly generated public key from the target service. The initiator fills in appropriate values for the fields in the Proxy Certificate as described in Section 3.1 as well as a policy describing the rights it wishes to delegate.

5. The new Proxy Certificate is sent back over the secured channel to the target service, which places it into a file with the newly generated private key. This new Proxy Certificate is then available for use on the target service for applications running on the user's behalf.

While the host receiving the delegated Proxy Certificate may have a long-term key pair of its own (bound to an X.509 public key certificate that it used for authentication), this key pair is typically not reused for the delegated Proxy Certificate. The reason is that a given host may have multiple users delegating privileges to it that are intended to be bound to specific processes and not shared across processes. The generation of a new key pair for each process greatly simplifies the task of keeping privileges compartmentalized. This approach also allows a user to “revoke” the delegation by deleting the Proxy Certificate private key as described in [13]. We discuss the performance ramifications of this approach in Section 7.

5 Authorization Models for Proxy Certificates

Proxy Certificates have three obvious modes of integration with authorization systems: full delegation of rights from the issuer—in effect, impersonation; no delegation of rights from the issuer, solely using attribute assertions to grant privileges; and a restricted delegation of some subset of the issuer's rights to the Proxy Certificate bearer. In this section we describe our experiences with each of these three methods.

5.1 *Identity-based Authorization with Impersonation*

In our initial implementations, we used Proxy Certificates almost exclusively as impersonation credentials that granted the bearer the full rights of the issuer. This approach has the advantage of integrating easily with identity-based authorization systems since these systems can simply treat the bearer of such a Proxy Certificate as they would its issuer. However, such usage is not ideal from the point of view of trying to achieve least privilege delegation since it only supports the full delegation of the issuer's rights. Thus, we explored other methods.

5.2 *Proxy Certificates with Restricted Delegation*

Proxy Certificates can be created with policies that delegate only a subset of the issuer's rights to the Proxy Certificate bearer. While this form of usage is more in line with the goal of enabling least privilege delegation, the implementation becomes more complex. As we described in Section 3.1, Proxy Certificates do not mandate any particular delegation language for the issuer to express their delegation policy, but instead provide a framework for containing policy statements using a method of the issuer's choosing.

The primary complication is that the relying party accepting the restricted Proxy Certificate must both understand the semantics of the delegation policy used and be able to enforce the restrictions that it imposes. Since these policies often contain application-specific restrictions, it is difficult for a security library handling the authentication of the Proxy Certificate to know what restrictions the application understands and is capable of enforcing. Without assurance that the application (or some other part of the software stack) will handle the enforcement of the restrictions, the authentication library cannot safely accept a restricted Proxy Certificate.

We have attempted to solve this problem by extending the API between the application and the security libraries to allow the application to express to the security libraries its knowledge and ability to handle given restriction delegation policies. However, this approach is difficult in practice since it must be done on a per-application basis. For this reason, we have not used this form of Proxy Certificate authorization to a large degree.

5.3 Identity Creation with Additional Assertions

The third method of using Proxy Certificates in authorization systems is to have Proxy Certificates convey no rights to the bearer (i.e., a policy type of “independent” as described in Section 3.1) and then use attribute assertions to assign rights to the bearer. This method has the advantage that attributes may be granted to the bearer from a number of different sources and may be done so at times other than the creation of the Proxy Certificate.

However, there are two difficulties in implementation of this method that have slowed our adoption:

- *Lack of protocol support:* The TLS protocol [6] and implementation of OpenSSL [32] (before the latest, version 0.9.7) lack support for X.509 attribute certificates. Thus, every application protocol must be modified to include a means of transporting attribute certificates. (We do note that our recent move to a web service based protocol [39] may ease this burden.)
- *Lack of granularity in enforcement systems:* Many enforcement systems do not have the ability to enforce any policies with finer granularity than simple groups. Although there has been some work in finer-grained enforcement [25,33,12,21,22] these results are not yet portable across all applications and operating systems.

6 Proxy Certificate Implementations and Applications

Here we briefly describe our implementation of Proxy Certificates and some applications that use Proxy Certificates.

6.1 Implementation in Globus Toolkit's Grid Security Infrastructure

The Grid Security Infrastructure (GSI) implements Proxy Certificates to provide authentication and delegation capabilities for the Globus Toolkit. It allows application users to employ proxy certificates to authenticate to GSI-based services and to delegate Proxy Certificates to those services so that they may act on the user's behalf.

GSI is primarily intended to work with identity-based authorization systems and as such returns to the calling application an identity for the remote client. It is further intended to be used primarily with Proxy Certificates that have policies delegating the full set of their issuer's rights to their bearer. In this case it returns the subject name from the X.509 public key certificate that issued the original Proxy Certificate in the chain. As we describe in Section 6.4, GSI has also been used successfully with a combination of Proxy Certificates and attribute assertions. The use of GSI with restricted Proxy Certificates has been hampered by the issues described in Section 5.2.

GSI includes a GSS-API [24] library, which handles authentication and delegation using Proxy Certificates. This library is based heavily on the OpenSSL [32] library, an open source implementation of the SSL protocol. The library uses OpenSSL to provide protocol support, including message protection and basic X.509 path validation. It adds to OpenSSL custom code for handling Proxy Certificates in addition to normal X.509 public key certificates and performing delegation.

6.2 MyProxy: An Proxy Certificate Repository

MyProxy [31,26] is a credential repository service that enables credential mobility and also alleviates the burden on users of managing and protecting files containing long-term secrets (i.e., private keys). We describe MyProxy briefly here, directing readers interested in more information to the references.

MyProxy is similar in function to a traditional credential repository as defined in the IETF SACRED working group [18]. However, by using Proxy Certificates it can operate without long-term private keys ever leaving the MyProxy service. MyProxy allows a user to establish a protected channel to the MyProxy service using SSL (without a client-side certificate), to authenticate over that channel from a remote system using, for example, a username and pass phrase, and then obtain a Proxy Certificate bearing their privileges without having to carry their long-term public key certificate and private key around with them (a potentially error-prone and insecure process).

6.3 Use in other Applications

The GSI libraries have also found uses in common applications. For example, Proxy Certificates can be used as an alternative authentication mechanism in secure shell (SSH) [15], CVS [14], and FTP [1]. These and other applications use the GSS-API library from GSI to allow a user to authenticate to an appropriate GSI-enabled daemon using their Proxy Certificate. The GSI-enabled SSH application also allows the user to delegate a new Proxy Certificate so that other GSI-enabled applications can be used on the remote system.

6.4 Proxy Certificates as Attribute Assertion Carriers

Combining public-key certificates with attribute assertions allow for the reuse of a single PKI across multiple application domains. In such a scenario, the PKI is used as a identity provider and all applications or domain specific privilege information (e.g., group memberships, clearance level, citizenship) is conveyed by separate attribute authorities.

However, as we mention in Section 5.3, many security protocols do not offer support for conveying attribute assertions. For example, the TLS protocol does not allow for attribute certificates in the set of provided client credentials. Thus, each application protocol must be modified to accommodate attribute assertions.

One way to circumvent this problem is by way of Proxy Certificates: when creating a Proxy Certificate, the proxy certificate issuer has the opportunity to add additional information to the proxy certificate by way of certificate extensions (in addition to the PCI extension described in Section 3.1). Several Grid projects use this technique to bundle application-specific attributes dynamically in the Proxy Certificate. The Community Authorization Service (CAS) [33,12] makes use of SAML authorization

decisions [34] to assert that the identity may perform (a group of) actions on (a group of) objects. The VO Membership Service (VOMS) [38] is a role-based authorization system that uses X.509 attribute certificates to assert a user's group membership(s), role(s), and capabilities. PRIMA [25] is a similar system that uses X.509 attribute certificates containing XACML [7] statements to assert a user's capabilities.

7 Performance and Security Issues

The expensive part of a Proxy Certificate creation is generating the new key pair. In this section, we only consider RSA key pairs due to lack of support in commonly used open source software stacks for alternatives, such as elliptic curve cryptography (ECC) [27].

Generating an RSA public key pair involves finding a pair of suitable prime numbers, which is a non-trivial amount of work that furthermore scales exponentially with the key length. Table 2 shows timings for key pair generation on a 2.8GHz Pentium 4 processor using the OpenSSL 0.9.7 library. We measure system CPU time and give averages over 100 keys.

Size (bits)	Time (seconds)
512	0.040
768	0.094
1024	0.176
1536	0.415
2048	1.348

Table 2: Key generation times for RSA key pairs

Unfortunately, use of specialized hardware such as cryptographic accelerators does not help these timings much, as such hardware is built with the assumption that RSA key generation occurs seldom and thus is not a performance sensitive operation.

Consequently, key generation of normal key sizes may consume a substantial amount of CPU for hosts receiving delegated Proxy Certificates from multiple clients. It is tempting to use smaller key sizes since the lifetime of a Proxy Certificate key pair is comparably short. (Indeed, the 3.0 release of the Globus Toolkit does just this.) While a smaller key size may yet meet the targets for complexity necessary to make brute force attacks infeasible within the short lifetime of the key pair, one must remember the cascading effects on the context in which such a key is used. For example, private data transferred during an ftp connection will typically remain sensitive long after the transfer is completed, and if an eavesdropper records the whole ftp transfer they have a longer period of time than the life of the key pair during which they may attack the protection it provided.

Thus, we note that Proxy Certificate generation comes with a non-negligible penalty in server-side key generation. Currently this means that services must take appropriate precautions when accepting Proxy Certificate delegations to prevent denial of service

attacks. At the time of writing, the development of solutions that mitigate this problem is left as future work.

8 Related work

A number of schemes offer delegation in a similar manner to Proxy Certificates. We discuss a few of these schemes here and compare them to our Proxy Certificate work.

8.1 Kerberos V5

The Kerberos Network Authentication Protocol [23,29] is a widely used authentication system based on conventional (shared secret key) cryptography. It provides support for single sign-on via creation of “Ticket Granting Tickets” (TGTs), and support for delegation of rights via “forwardable” and “proxyable” tickets. The initial use of proxy credentials in Kerberos was described by Neuman [30], who also described restricted proxy credentials and proposed several uses for them, including cascaded delegation (using a proxy credential that contains restrictions to generate a new proxy with greater restrictions), authorization servers (servers that grant restricted proxy credentials based on a database of authorization information), and group servers (servers that grant restricted proxy credentials that convey rights to assert membership in groups).

From the perspective of a user, applications using Kerberos 5 are similar to applications using X.509 Proxy Certificates. The features of Kerberos 5 tickets formed the basis of many of the ideas surrounding X.509 Proxy Certificates. For example, the local creation of a short-lived Proxy Certificate can be used to provide single sign-on in an X.509 PKI based system, just as creation of short-lived TGT allows for single sign-on in a Kerberos based system. And a Proxy Certificate can be delegated just as a forwardable ticket can be forwarded. Proxy Certificate and Kerberos also share the common method of protecting a TGT and protecting the private key of a Proxy Certificate by using local filesystem permissions.

The major difference between Kerberos TGTs and X.509 Proxy Certificates is that creation and delegation of a TGT requires the involvement of a third party (the Kerberos Domain Controller), while Proxy Certificates can be unilaterally created by their issuers without the active involvement of a third party.

8.2 X.509 Attribute Certificates

An X.509 attribute certificate (AC) [8] can be used to grant to a particular identity some attribute such as a role, clearance level, or alternative identity such as “charging identity” or “audit identity.” Authorization decisions can then be made by combining information from the identity itself with signed attribute certificates providing binding of that identity to attributes. Attribute certificates can either be issued by a trusted entity specific to the issuance of attributes, known as an attribute authority, or by end entities delegating their own privileges.

In the case of an attribute authority, this method works equally well with attributes certificates bound to public key certificates or Proxy Certificates. For example, Proxy Certificates can be used to delegate the issuer’s identity to various other parties who can claim attributes of the issuer. An AC could also be bound directly to a particular Proxy Certificate using the unique subject name from the Proxy Certificate.

The uses of ACs that are granted directly by end entities overlap considerably with the uses of Proxy Certificates. However, this AC based solution to delegation has some disadvantages as compared to the Proxy Certificate based solution:

- A similar modification to the validation framework, as in the Proxy Certificate RFC and described in Section 3.2, is needed in order to allow ACs to be signed by end entities.
- Identifying short-lived, dynamically created identities as described in Section 2.2, remains a non-resolved problem.
- All protocols, authentication code, and identity based authorization services must be modified to understand ACs.
- ACs must be created and signed by the long-term identity credentials of the end entity. This implies that the entity must know in advance which other identities may be involved in a particular task in order to generate the appropriate ACs. On the other hand, Proxy Certificates bearers can delegate privileges through the creation of new Proxy Certificates without interaction of the entity holding the long-term identity credentials.

We believe there are many unexplored tradeoffs between ACs and Proxy Certificates. Reasonable arguments can be made in favor of either an AC-based solution to delegation or a Proxy Certificate based solution to delegation. The approach to be taken in a given instance may depend on factors such as the software that it needs to be integrated into, the type of delegation required, and religion.

8.3 SPX

SPX [36] uses a structure entitled a “ticket” for delegation and single sign-on which is similar in purpose to Proxy Certificates. The two mechanisms share many common features: the SPX ticket is combined with a private key to provide a set of credentials to provide the means for authentication; the ticket and its private key are short-lived and normally stored in a file protected by file permissions; and the implementation uses the GSS-API as the application interface.

The main difference is that SPX defines its own format for the ticket and its own protocols for authentication. Proxy Certificates, being based on X.509 public key certificates, allow for a significant reuse of the existing protocols and software designed for those certificates.

Proxy Certificates also include the concept of a delegation policy (Section 3.1), which allows for arbitrary delegation of subsets of the issuers rights to the Proxy Certificate bearer. In contrast, SPX tickets only offer an impersonation mode.

8.4 Delegation in Digital's DSSA

Gasser and McDermott [13] describe a delegation scheme used in Digital's Distributed System Security Architecture (DSSA). This restricted public-key based delegation is similar to Proxy Certificates in that it allows for cascading delegation, has delegations bound to unique keys, and has similar motivations. The primary difference between Proxy Certificates and the DSSA work is our starting from X.509 public key certificates

in order to allow for maximum protocol and software reuse. It is also unclear to what extent the DSSA work was implemented.

8.5 Future XML Alternatives

Proxy Certificates offer a pragmatic approach to delegation of rights in a SSL- and X.509-dominated world. By basing Proxy Certificates on the well established X.509 certificates, the Proxy Certificates chains are easily exchanged in the SSL authentication protocol. Furthermore, by embedding the delegation policy statements inside of the Proxy Certificate, these delegation directives are exchanged as part of the SSL authentication process

At this time, we appear to be moving towards a web services dominated world. We envision that pure XML-based alternatives to SSL/TLS will be invented for authentication and key exchange based on new and emerging specifications and standards, such as XML-Signature, XML-Encryption, WS-Trust, WS-SecureConversation, etc. We expect these new standards to be more authentication mechanism agnostic and supporting alternatives to X.509, such as PGP, SPKI, bare keys, etc. Furthermore, these protocols are also expected to be able to communicate attribute and authorization assertions transparently without requiring modification of the application protocol. Some of our initial work in this area is described in [39].

We are currently investigating these XML-based technologies as alternatives or enhancements to Proxy Certificates. For example, the equivalent functionality of a Proxy Certificate could be achieved through a fine-grained SAML [34] authorization assertion expressed or an XACML policy statement that empowers a bare key. The generation of this key and the issuing of this authorization assertion could follow the same procedure and pattern as we use for Proxy Certificates.

9 Summary

Standard X.509 identity and attribute certificates allow for the static assignment of identities and rights. However, some environments require that end entities be able to delegate and create identities quickly. We have described Proxy Certificates, a standard mechanism for dynamic delegation and identity creation in public key infrastructures. Proxy certificates are based on X.509 public key certificates in order to allow for significant reuse of protocols and open source software. Our Grid Security Infrastructure (GSI) implementation of Proxy Certificates exploits these opportunities for reuse to provide a widely used implementation of Proxy Certificate mechanisms. A number of applications and widespread deployment demonstrate the viability of Proxy Certificate mechanisms.

10 Acknowledgements

We are pleased to acknowledge significant contributions to the Proxy Certificate RFC by David Chadwick, Doug Engert, Jim Schaad, and Mary Thompson. We are also grateful to numerous colleagues for discussions regarding Proxy Certificates, in particular: Carlisle Adams, Joe Bester, Randy Butler, Keith Jackson, Steve Hanna, Russ Housley, Stephen Kent, Bill Johnston, Marty Humphrey, Sam Lang, Ellen McDermott, Clifford Neuman, and Gene Tsudik. Doug Engert coded the initial prototype implementation of Proxy

Certificates in GSI. Sam Meder, Jarek Gawor and Sam Lang coded the current implementations. We also thank Jim Basney and the anonymous members of the program committee for reviewing and commenting on early versions of this paper.

“Globus Toolkit” is a registered trademark of the University of Chicago.

This work was supported in part by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under contracts W-31-109-Eng-38, DE-AC03-76SF0098, DE-FC03-99ER25397 and No. 53-4540-0080.

11 References

1. Allcock, B., et. al., Data Management and Transfer in High Performance Computational Grid Environments. *Parallel Computing Journal*, Vol. 28 (5), May 2002, pp. 749-771.
2. Butler, R., Engert, D. Foster, I. , Kesselman, C. , Tuecke, S. , Volmer, J. , and Welch, V. A National-Scale Authentication Infrastructure. *IEEE Computer*, 33(12):60-66, 2000.
3. Beiriger, J., Johnson, W., Bivens, H., Humphreys, S. and Rhea, R., Constructing the ASCI Grid. *In Proc. 9th IEEE Symposium on High Performance Distributed Computing*, 2000, IEEE Press.
4. Brunett, S., Czajkowski, K., Fitzgerald, S., Foster, I., Johnson, A., Kesselman, C., Leigh, J. and Tuecke, S., Application Experiences with the Globus Toolkit. *In Proc. 7th IEEE Symp. on High Performance Distributed Computing*, 1998, IEEE Press, 81-89.
5. CCITT Recommendation, X.509: The Directory – Authentication Framework. 1988.
6. Dierks, T. and Allen, C., The TLS Protocol Version 1.0, *RFC 2246*, IETF, 1999.
7. eXtensible Access Control Markup Language (XACML) 1.0 Specification, OASIS, February 2003.
8. Farrell, S. and Housley,R., An Internet Attribute Certificate Profile for Authorization, *RFC 3281*, IETF, April 2002.
9. Foster, I., Kesselman, C., Tsudik, G. and Tuecke, S. A Security Architecture for Computational Grids. *ACM Conference on Computers and Security*, 1998, 83-91
10. Foster, I. and Kesselman, C. Computational Grids. Foster, I. and Kesselman, C. eds. *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999, 2-48.
11. Foster, I. and Kesselman, C. Globus: A Toolkit-Based Grid Architecture. Foster, I. and Kesselman, C. eds. *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999, 259-278.
12. Foster, I., Kesselman, C. , Pearlman, L., Tuecke, S., Welch, V. , The Community Authorization Service: Status and future. *Proceedings of the International Conference on Computing in High Energy Physics - CHEP 2003*.
13. Gasser, M. and McDermott, E., An Architecture for Practical Delegation in a Distributed System. *Proc. 1990 IEEE Symposium on Research in Security and Privacy*, 1990, IEEE Press, 20-30.
14. gridCVS, <http://www.globus.org/gridcvs/>, 2002.
15. GSI-Enabled OpenSSH, <http://grid.ncsa.uiuc.edu/ssh/>, 2004.
16. Housley, R., Polk, W., Ford, W., and Solo, D., Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. *RFC 3280*, IETF, April 2002.
17. IETF Public-Key Infrastructure (X.509) (pkix) working group. <http://www.ietf.org/html.charters/pkix-charter.html>, January 2004.
18. IETF Securely Available Credentials (sacred) working group. <http://www.ietf.org/html.charters/sacred-charter.html>, 2003.
19. Johnston, W.E., Gannon, D. and Nitzberg, B., Grids as Production Computing Environments: The Engineering Aspects of NASA's Information Power Grid. *In Proc. 8th IEEE Symposium on High Performance Distributed Computing*, 1999, IEEE Press.

20. Kaliski, B., PKCS #10: Certification Request Syntax v1.5, *RFC 2314*, October 1997.
21. Keahey, K., Welch, V., Lang, S., Liu, B., Meder, S., Fine-Grain Authorization Policies in the GRID: Design and Implementation. *1st International Workshop on Middleware for Grid Computing*, 2003.
22. Keahey, K., and Welch, V. Fine-Grain Authorization for Resource Management in the Grid Environment. *Proceedings of Grid2002 Workshop*, 2002.
23. Kohl, J., and Neuman, C., The Kerberos Network Authentication Service (V5), *RFC 1510*, IETF, 1993.
24. Linn, J. Generic Security Service Application Program Interface, Version 2. *RFC 2078*, 1997.
25. Lorch, M., Adams, D., Kafura, D., Koneni, M., Rathi, A., and Shah, S., The PRIMA System for Privilege Management, Authorization and Enforcement in Grid Environments, *4th Int. Workshop on Grid Computing - Grid 2003*, 17 November 2003 in Phoenix, AR, USA.
26. Lorch, M., Basney, J., and Kafura, D., A Hardware-secured Credential Repository for Grid PKIs, *4th IEEE/ACM International Symposium on Cluster Computing and the Grid*, Chicago, Illinois, April 19-22, 2004 (to appear).
27. Menezs, A., Elliptic Curve Public Key Cryptosystems, *Kluwer Academic Publishers*, 1993
28. Myers, M., et. al., X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP, *RFC 2560*, IETF, June 1999.
29. Neuman, B. C. and Ts'o, T. Kerberos: An Authentication Service for Computer Networks. *IEEE Communications Magazine*, 32 (9). 33-88. 1994.
30. Neuman, B.C. Proxy-Based Authorization and Accounting for Distributed Systems. *In Proceedings of the 13th International Conference on Distributed Computing Systems*, pages 283-291, May, 1993.
31. Novotny, J., Tuecke, S., and Welch, V., An Online Credential Repository for the Grid: MyProxy. *Proceedings of the Tenth International Symposium on High Performance Distributed Computing (HPDC-10)*, IEEE Press, August 2001.
32. OpenSSL, <http://www.openssl.org>, 2002.
33. Pearlman, L., Welch, V., Foster, I., Kesselman, C. and Tuecke, S., A Community Authorization Service for Group Collaboration. *IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*, 2002.
34. Security Assertion Markup Language (SAML) 1.1 Specification, OASIS, November 2003.
35. Stevens, R., Woodward, P., DeFanti, T. and Catlett, C. From the I-WAY to the National Technology Grid. *Communications of the ACM*, 40(11):50-61. 1997.
36. Tardo, J.J. and Alagappan, K., SPX: global authentication using public key certificates, *Proceedings., 1991 IEEE Computer Society Symposium on Research in Security and Privacy*, Vol., Iss., 20-22 May 1991. Pages:232-244
37. Tuecke, S., Welch, V. Engert, D., Thompson, M., and Pearlman, L., Internet X.509 Public Key Infrastructure Proxy Certificate Profile, *draft-ietf-pkix-proxy-10 (work in progress)*, IETF, 2003.
38. VOMS Architecture v1.1, http://grid-auth.infn.it/docs/VOMS-v1_1.pdf, May 2002
39. Welch, V., et. al. Security for Grid Services, *Twelfth International Symposium on High Performance Distributed Computing (HPDC-12)*, IEEE Press, to appear June 2003.



the globus alliance

www.globus.org

X.509 Proxy Certificates for Dynamic Delegation

Ian Foster, Jarek Gawor, Carl Kesselman,
Sam Meder, Olle Mulmo, Laura Perlman,
Frank Siebenlist, Steven Tuecke,
Von Welch

(Presenter: vwelch@ncsa.uiuc.edu)

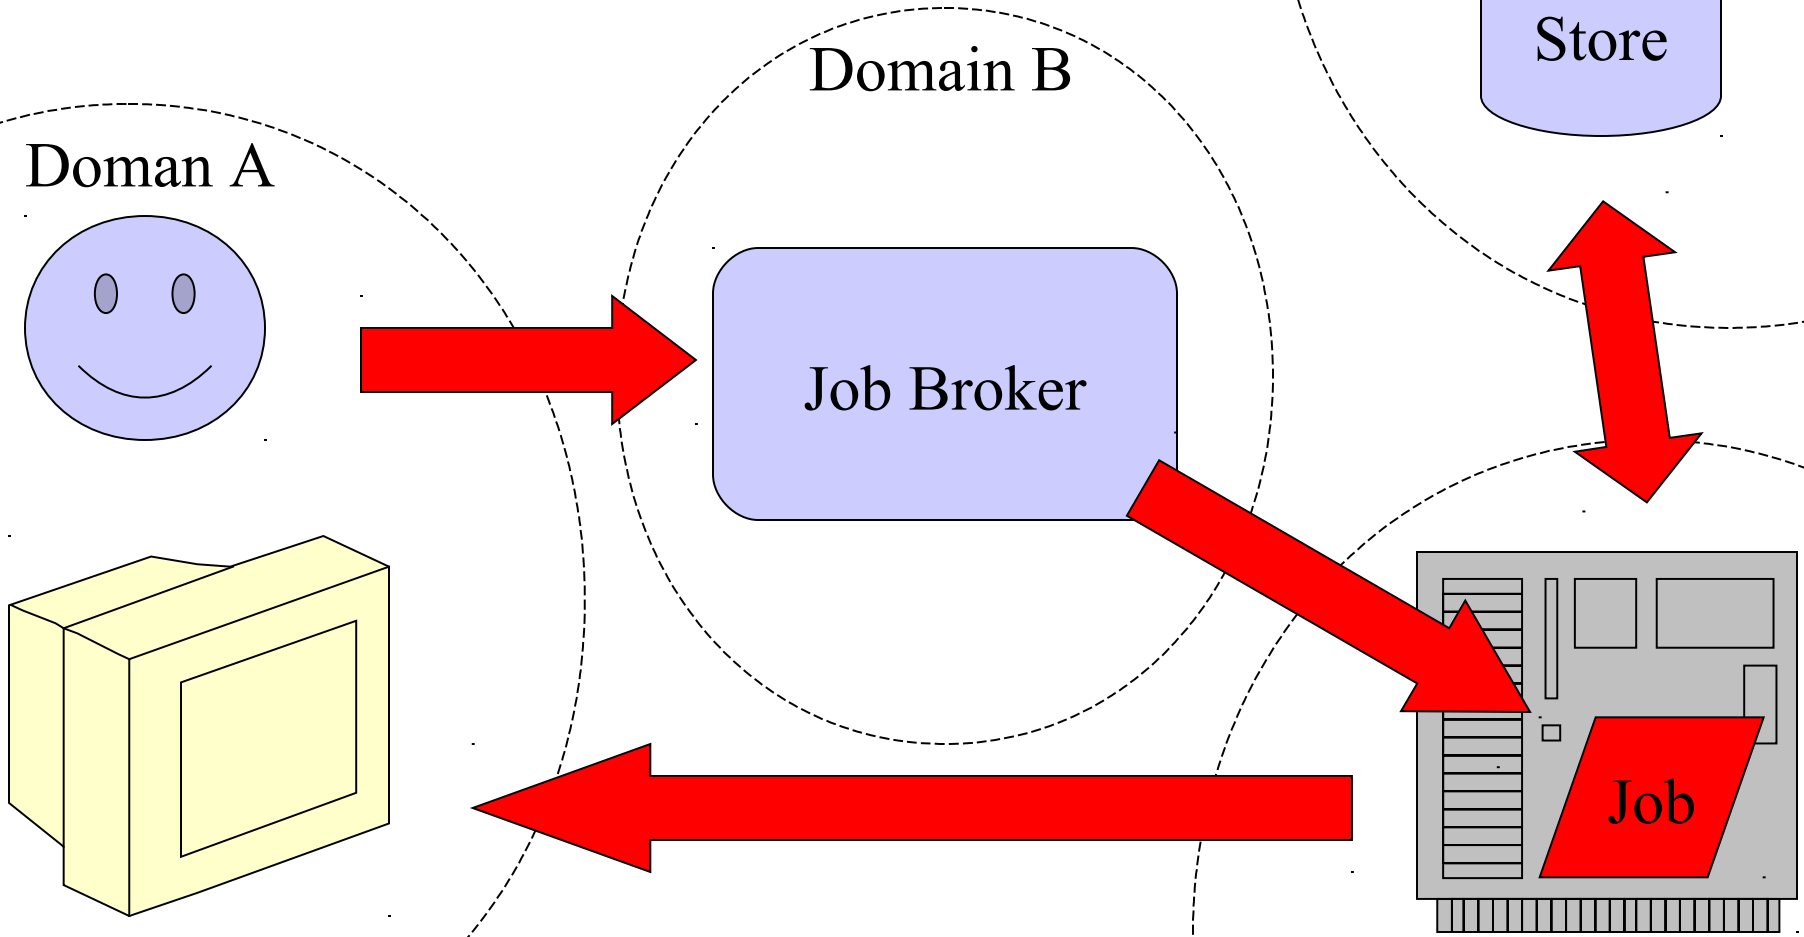


Outline

- Problem Statement, Motivations, Approach
- Proxy Certificate Solution
 - What are they?
 - What can they do?
- Status: Standardization, Implementation, Deployment



Use Case





Motivation

- **Dynamic Delegation**
 - Run-time decision on who and what
 - Support late binding of jobs to resources
- **Dynamic Entities**
 - Entities (e.g. Jobs) created at same time
- **Single Sign On**
 - Avoid repeated manual authentication
- **Easy (user-driven) cross-domain use**



Approach

- Start with PKI
 - Aids cross-domain trust issues since trust relationships can be set up by individual
- Build off of existing standards
 - Needs to be easily understood by security folks at many sites
- Ease of implementation
 - Use with existing PKI libraries as much as possible
 - Start with identity-based authz systems



Our solution: Proxy Certificates

- Allow users to delegate on the fly by granting other entities the right to use their name
- Prototypes in '98
- Standardized in IETF/PKIX 2004
- Fully implemented, deployed and widely used



Proxy Certificates

- Same format as X.509 Public Key Identify Certificate, but signed by user (or another proxy certificate)
- Name scoped to issuer's name
- Support restricted delegation from issuer to bearer
- Includes critical extension to identify as Proxy and express delegation



Issuer/ Signer	A certification authority	A public key certificate or another Proxy Certificate
Name	Any as allowed by issuer's policy	Unique, scoped to namespace defined by issuer's name
Delegation from Issuer	None	Allows for arbitrary delegation policies
Key pairs	Uses unique key pair	Uses unique key pair



ProxyCertInfo Extension

- Critical X.509 Extension
- Identifies a certificate as a Proxy Cert
- Allows issuer to express delegation intentions



ProxyCertInfo Delegation Policy

- Does not specify any method of expression
 - No language will be right for everyone all the time
- Instead OID to identify language and language-specific field
 - Any language can be used as long as understood by relying party
- Two methods defined: All and none

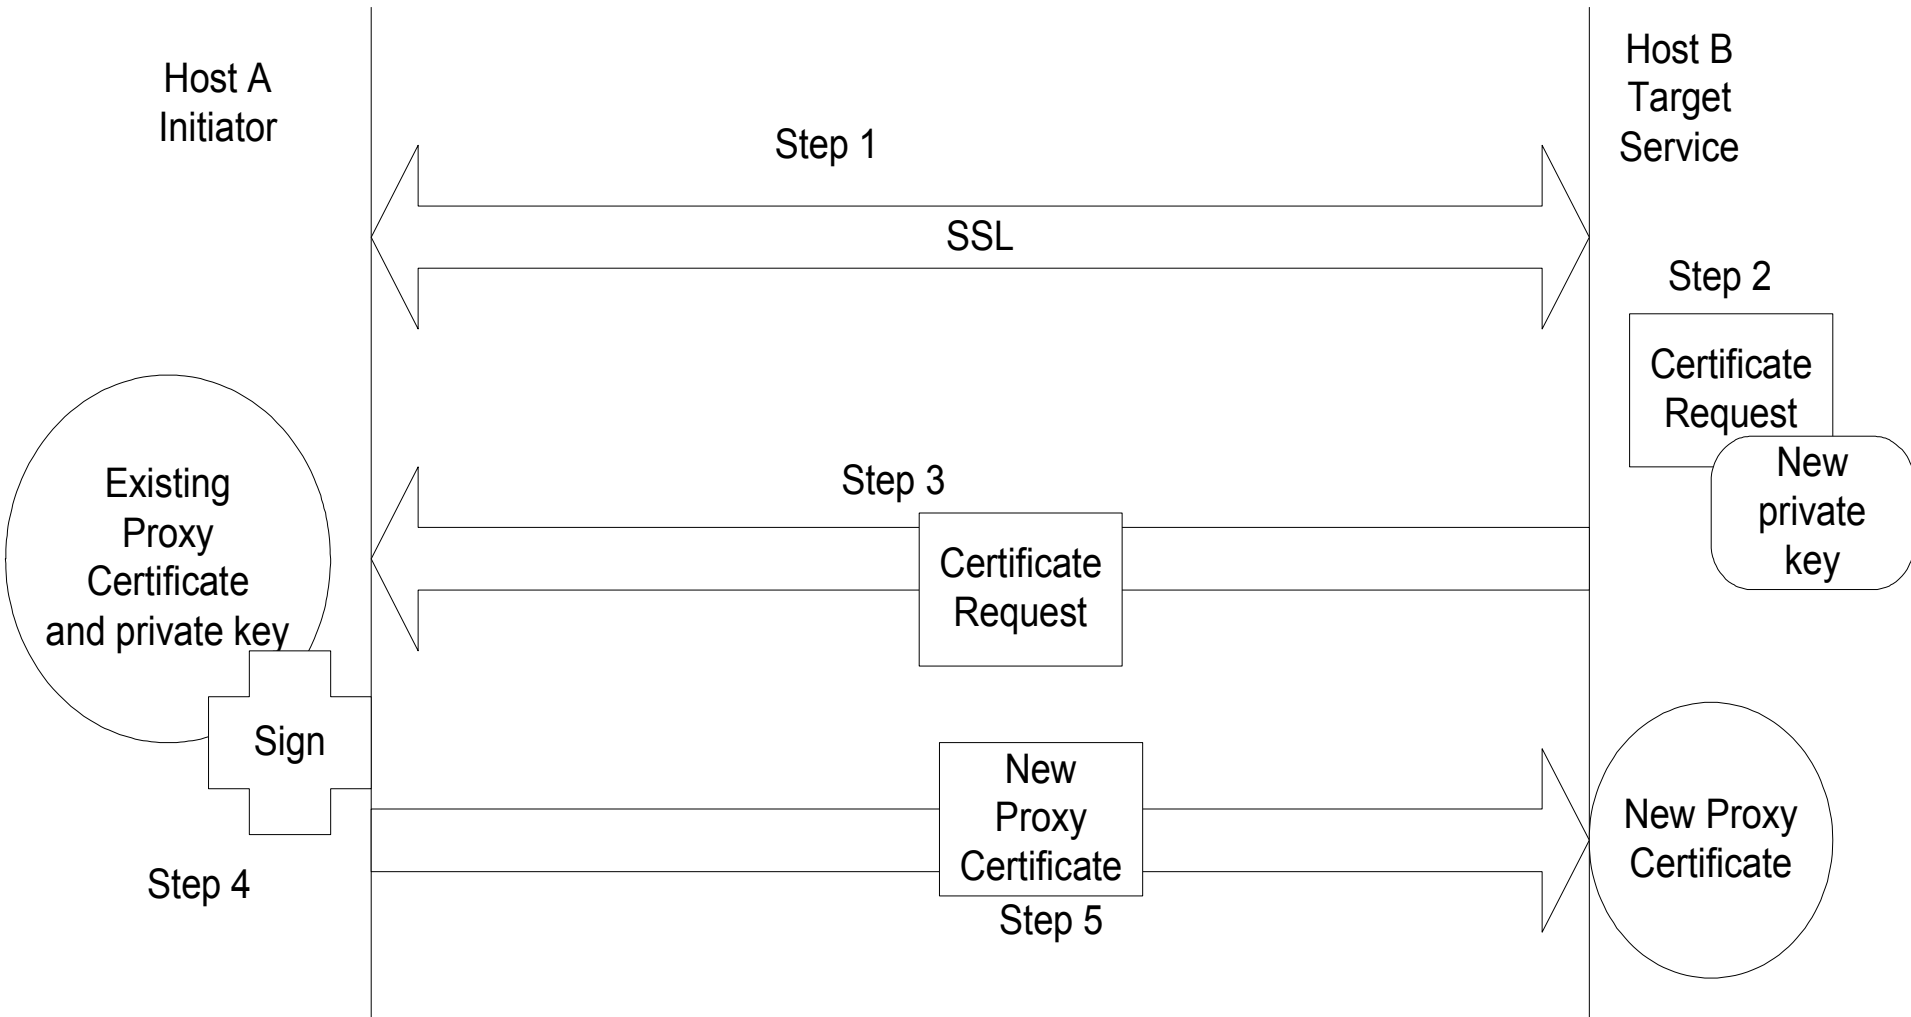


Single Sign On

- User creates key pair locally
- Signs new public key with identity private key
- Gives short life span
 - E.g. 8 hours
- Probably all rights
- Allows for weak (filesystem) protection of private key and easy use



Delegation





Performance and Security Issues

- Proxy generate requires key pair generation
- Those accepting delegation must take care to prevent DoS
 - Validate delegation request before generating key pair



Authorization Methods

- All rights/impersonation
 - Works great if you don't mind ignoring least privilege
- Delegation with restrictions
 - Issue: How does authentication mechanisms know restrictions will be enforced?
- Identity from Proxy Certificate plus additional assertions to grant rights



Standardization Status

- Proxy certificates have passed PKIX and IETF last calls
- Awaiting editorial process to become RFC
- Latest version is draft-ietf-pkix-proxy-10:
 - <http://www.ietf.org/internet-drafts/draft-ietf-pkix-proxy-10.txt>
 - Defines specifics of Proxy certificate creation and path validation



Implementation

- Fully implemented in Globus Toolkit's Grid Security Infrastructure (GSI)
 - www.globus.org/security/
- Build on OpenSSL
 - Changes are additions to handle Proxy Cert path validation as error handlers to normal path validation
- Similar Java implementation
- GSSAPI-based library
 - Also integrated with SSH, FTP, CVS



Deployment

- Many CAs issuing certificates for use with Proxy certificates for production Grids around the world
 - Master CA list at <http://www.gridpma.org/>
 - Two dozen plus CAs, including DOE, NSF, NASA
- Old Globus CA with 5k+ certs



Future Work

- One-time passwords/Two-factor authentication
 - Lot of recent attacks using keyboard sniffing
 - Service that hands out proxies authenticating with OTP
 - Poor man's hardware tokens
- Reasonable Restrictions
 - Where from? Intended use?
 - IP addresses too fragile (NAT, mobility, multi-homed)
 - Allow for late binding to resources
- Revocation
 - Even with short lifetime, interest in revocation



Summary

- Proxy Certificates are extension to X.509 identify certificates to allow for real-time delegation and naming
- Implemented with minimal changes to existing PKI libraries
- In production use in Grids world-wide
- Implementation available as part of Globus Toolkit (www.globus.org)



Acknowledgements

- DOE
 - SciDAC “Security for Group Collaboration”
- Many colleagues in Global Grid Forum and IETF for ideas and discussions
- Questions?

Panel: Controlled and Dynamic Delegation of Rights

Frank Siebenlist, ANL (moderator)
Von Welch, NCSA
Carl Ellison, Microsoft
Ravi Pandya, Microsoft
Kent Seamons, Brigham Young University

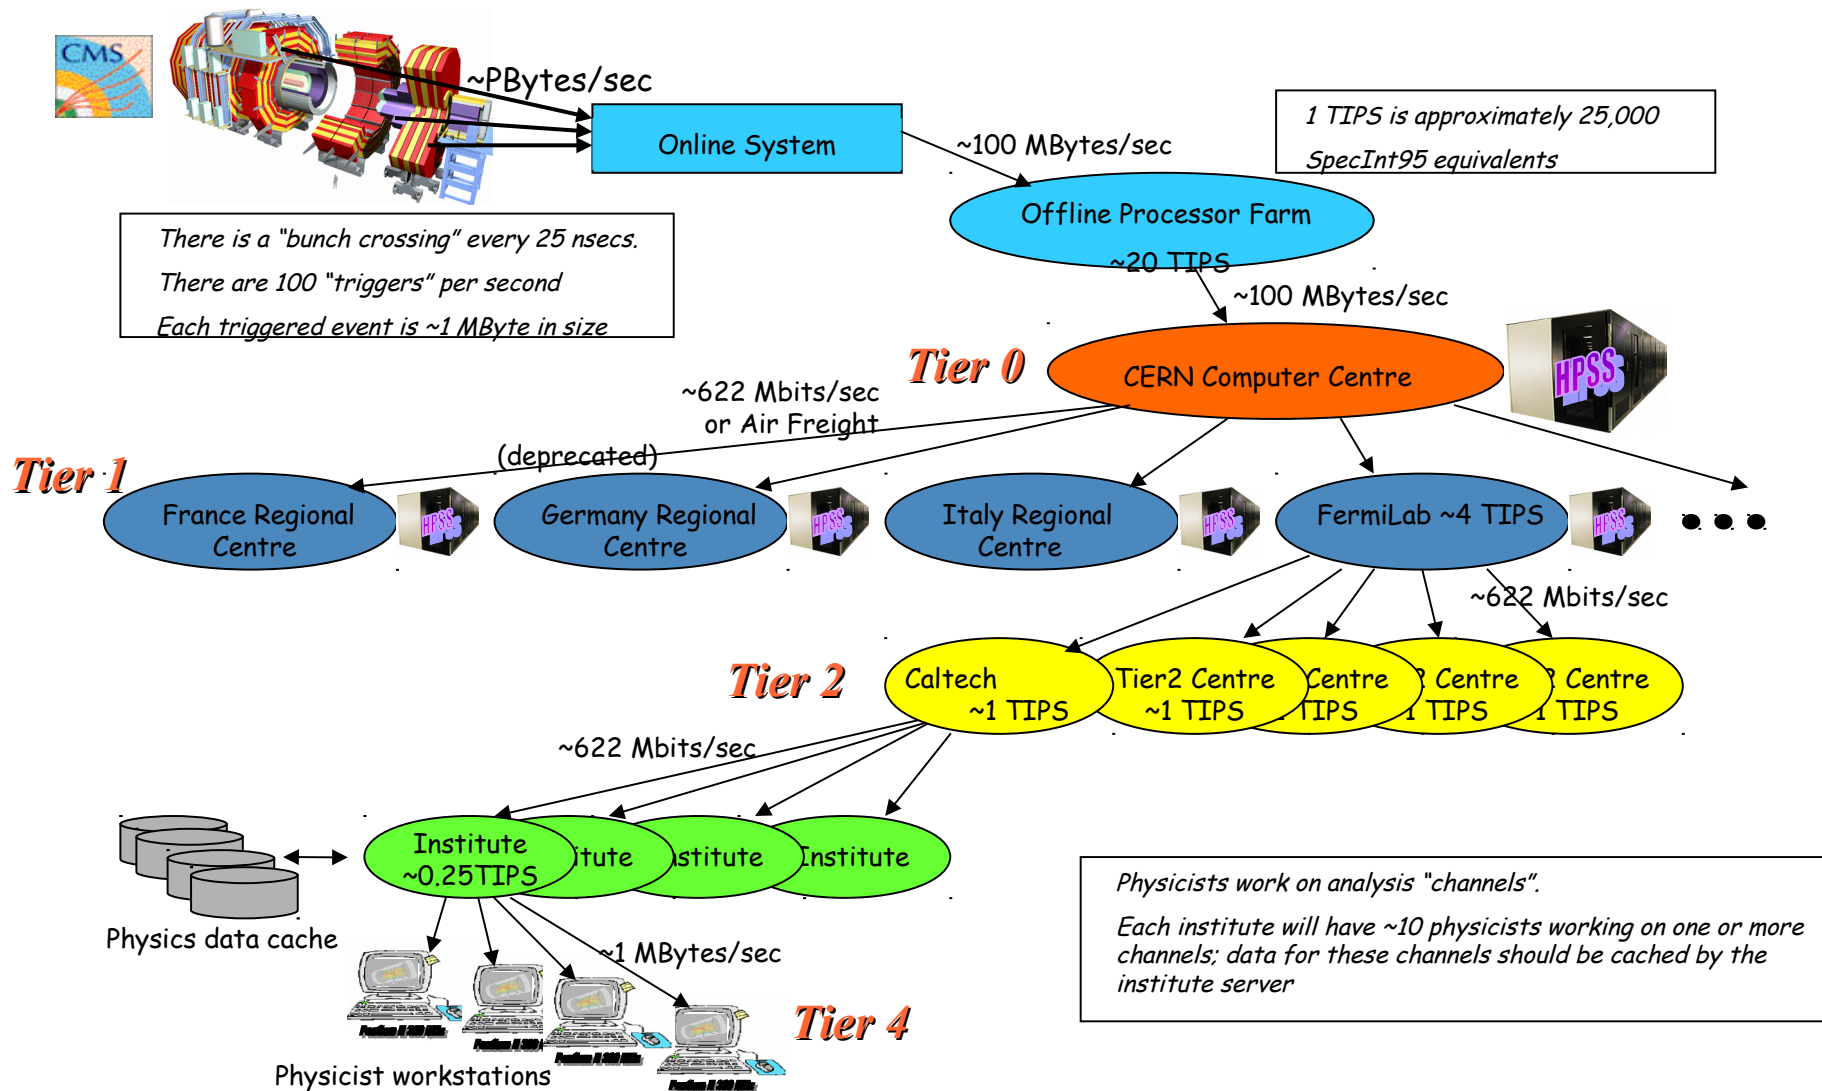
3rd Annual PKI R&D Workshop
April 12, 2004
NIST, Gaithersburg, MD

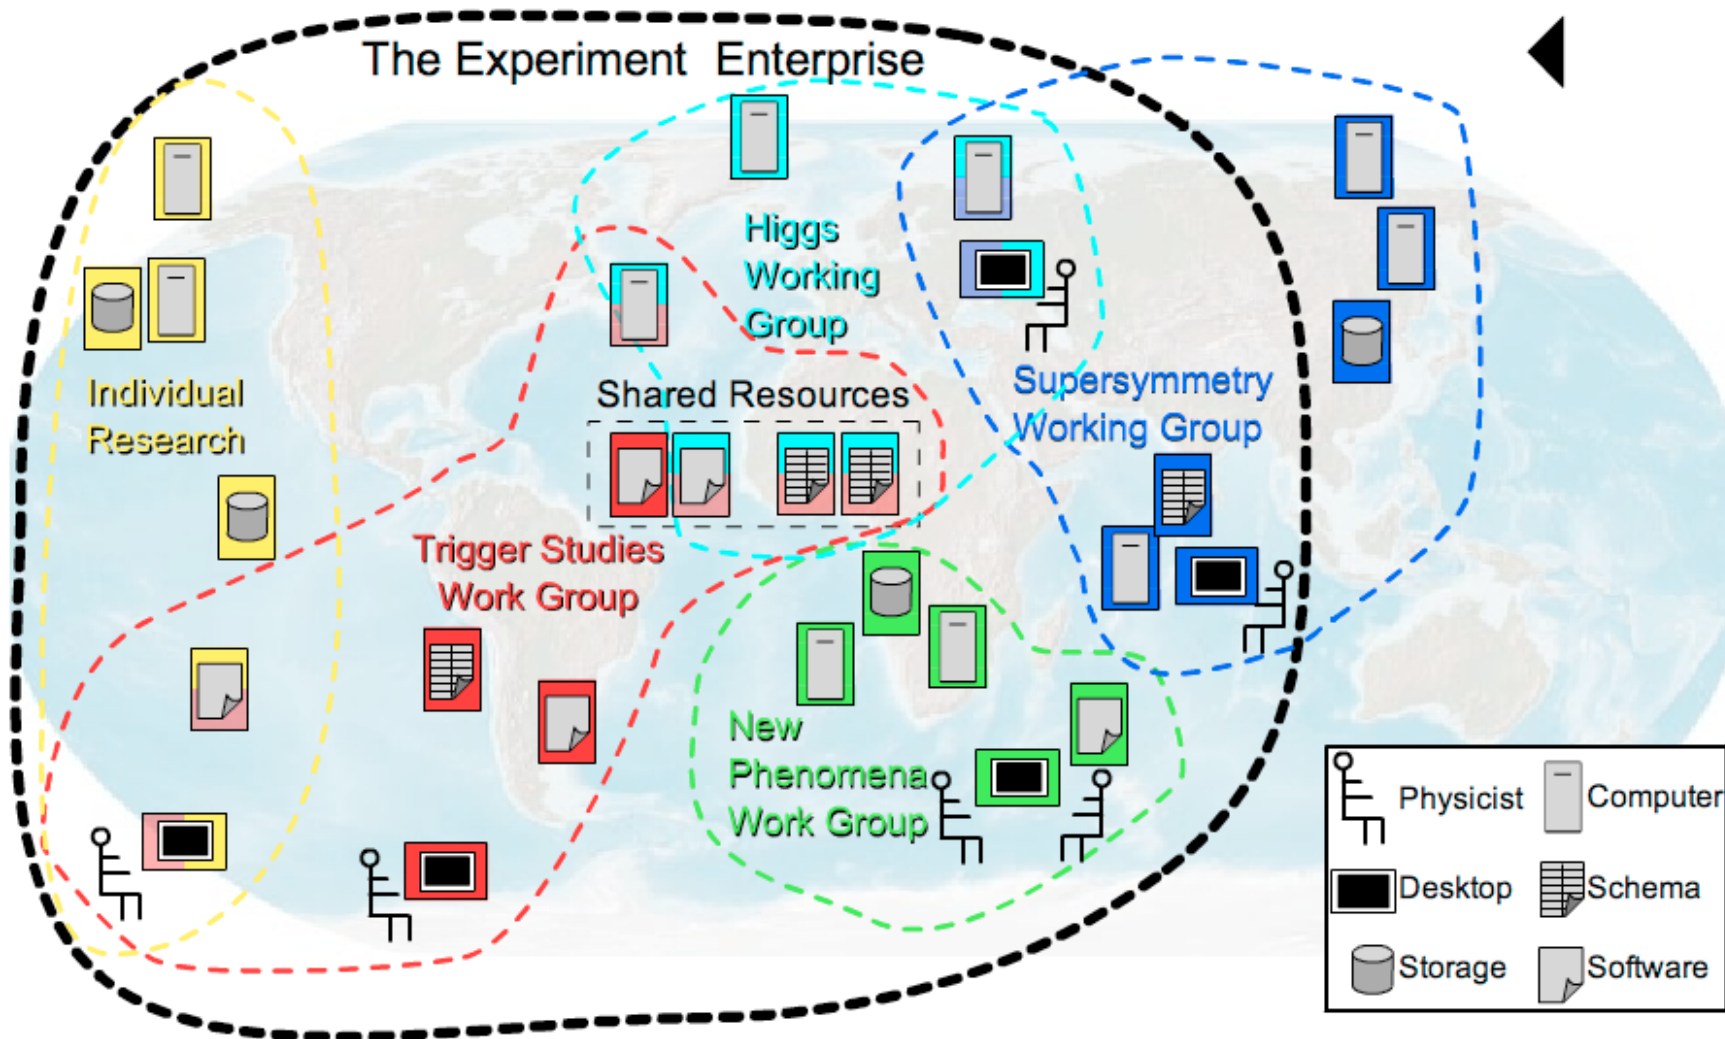
Outline

- X.509 Proxy-Certificates
 - ◆ Von Welch
- (Grid) Use Cases for Delegation of Rights
 - ◆ Frank Siebenlist
- SPKI (UPnP)
 - ◆ Carl Ellison
- XrML
 - ◆ Ravi Pandya
- TrustBuilder
 - ◆ Kent Seamons

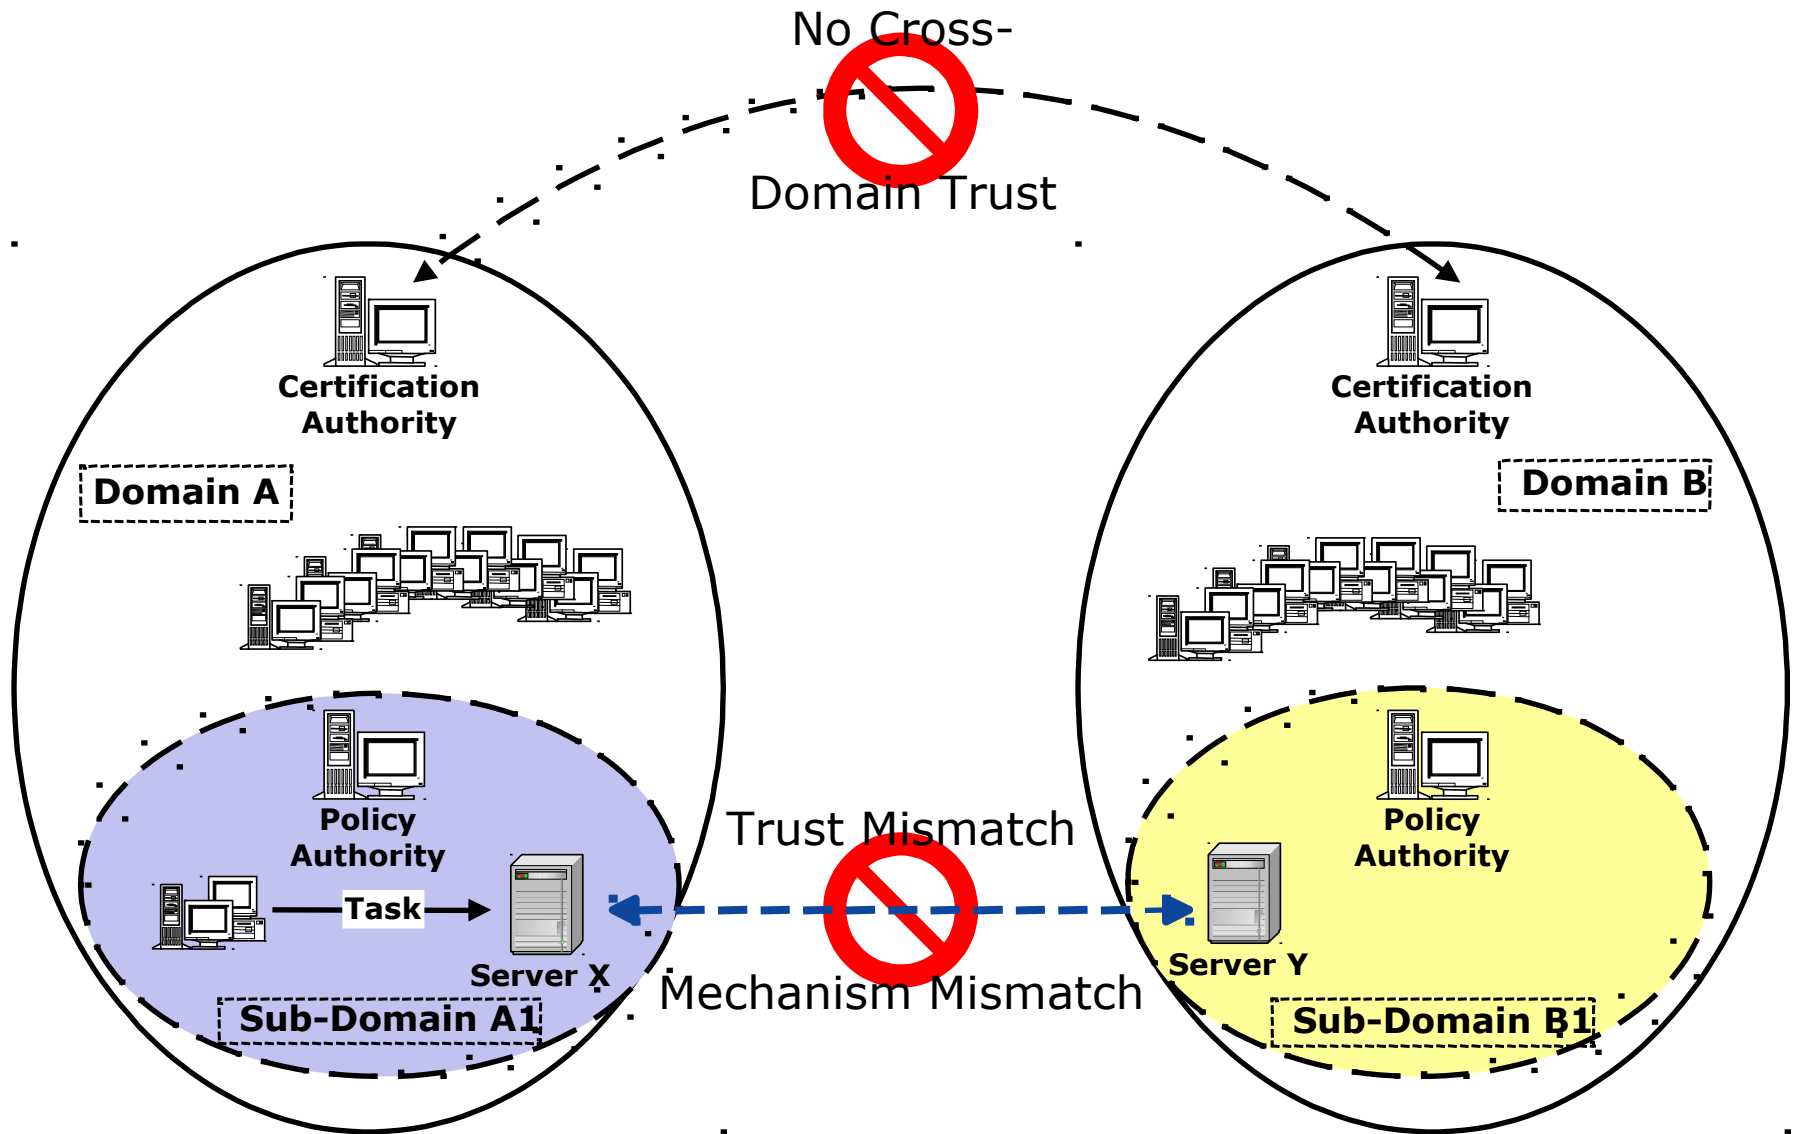
(Grid) Use Cases for Delegation of Rights

LHC Data Distribution

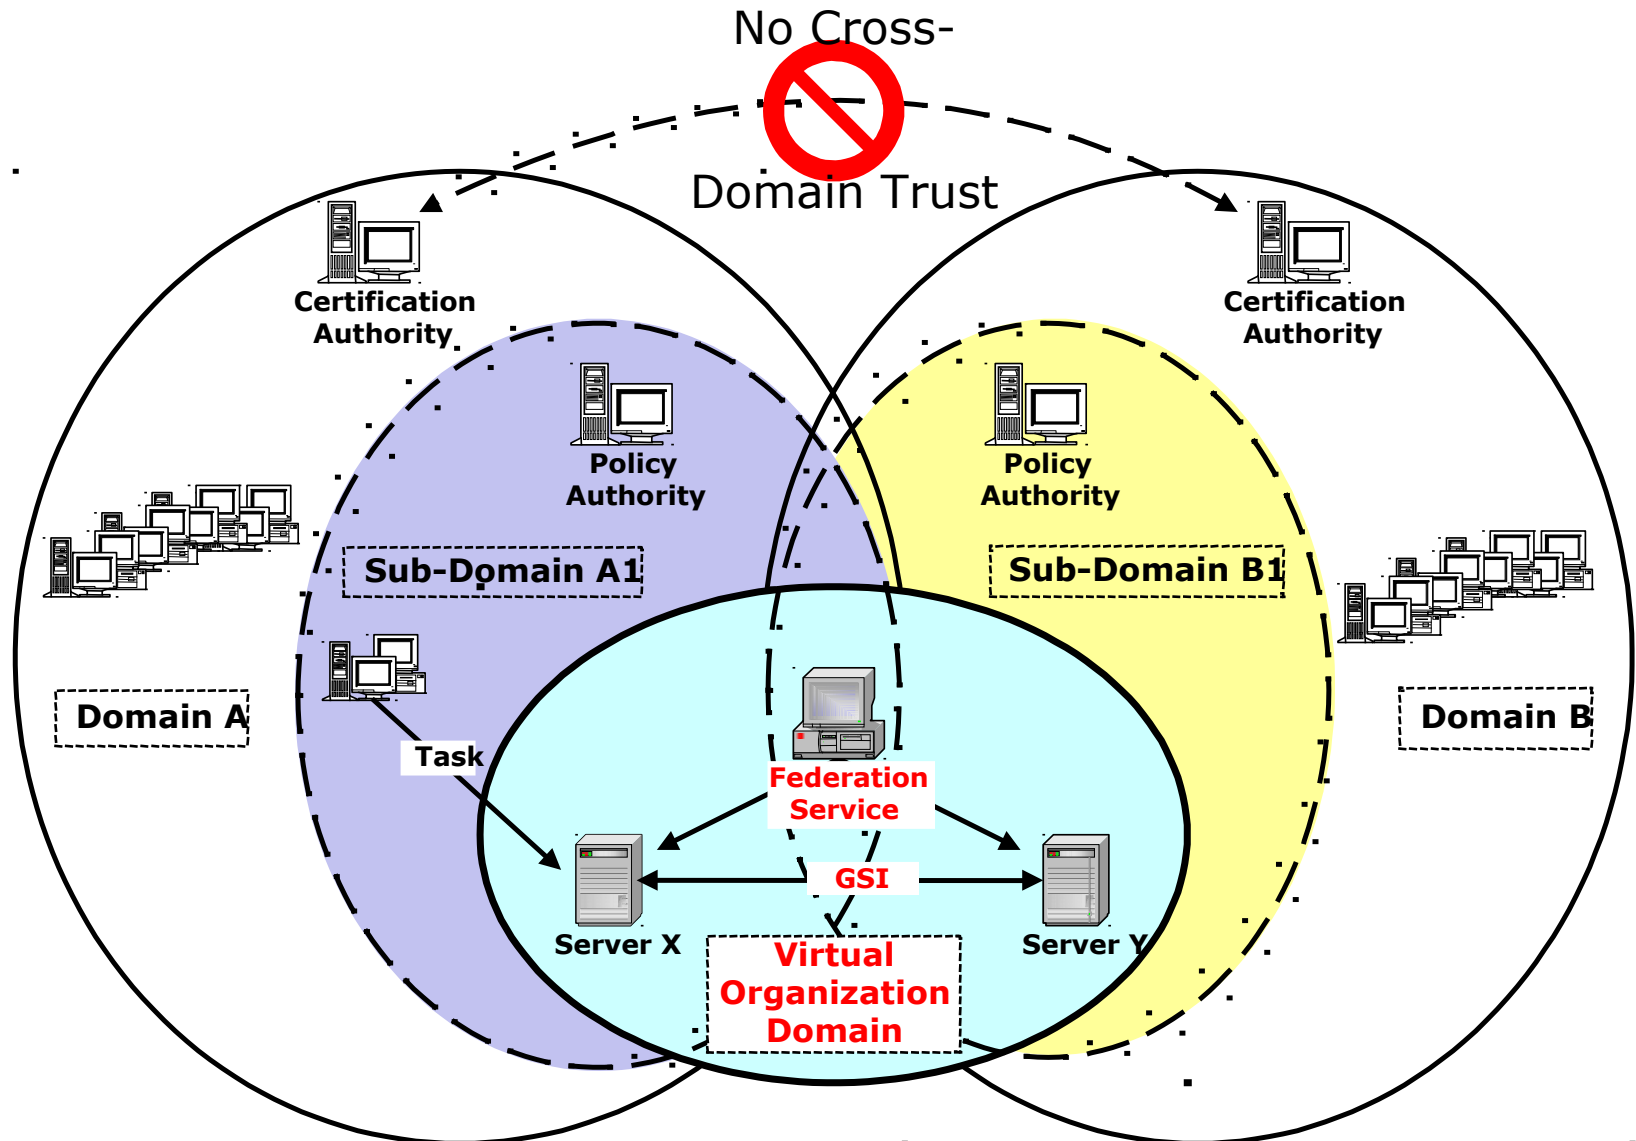




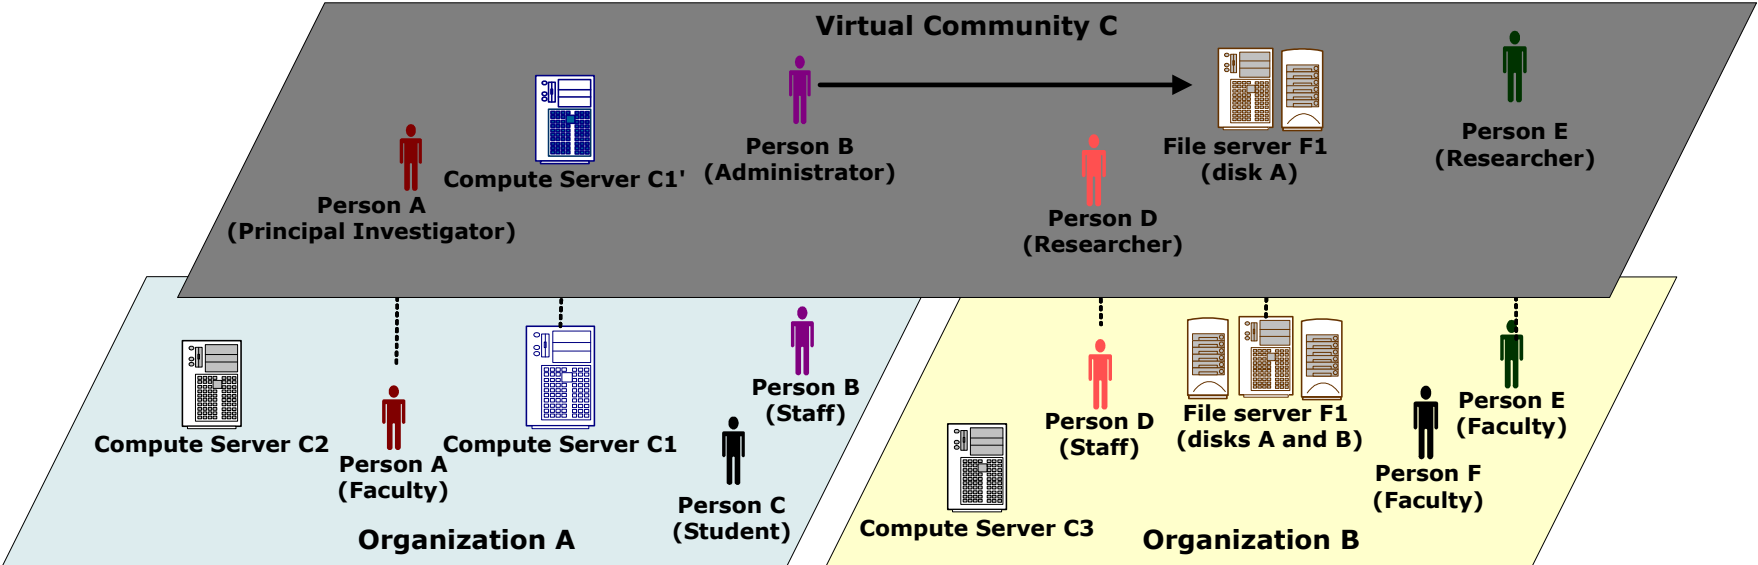
Multi-Institution Issues



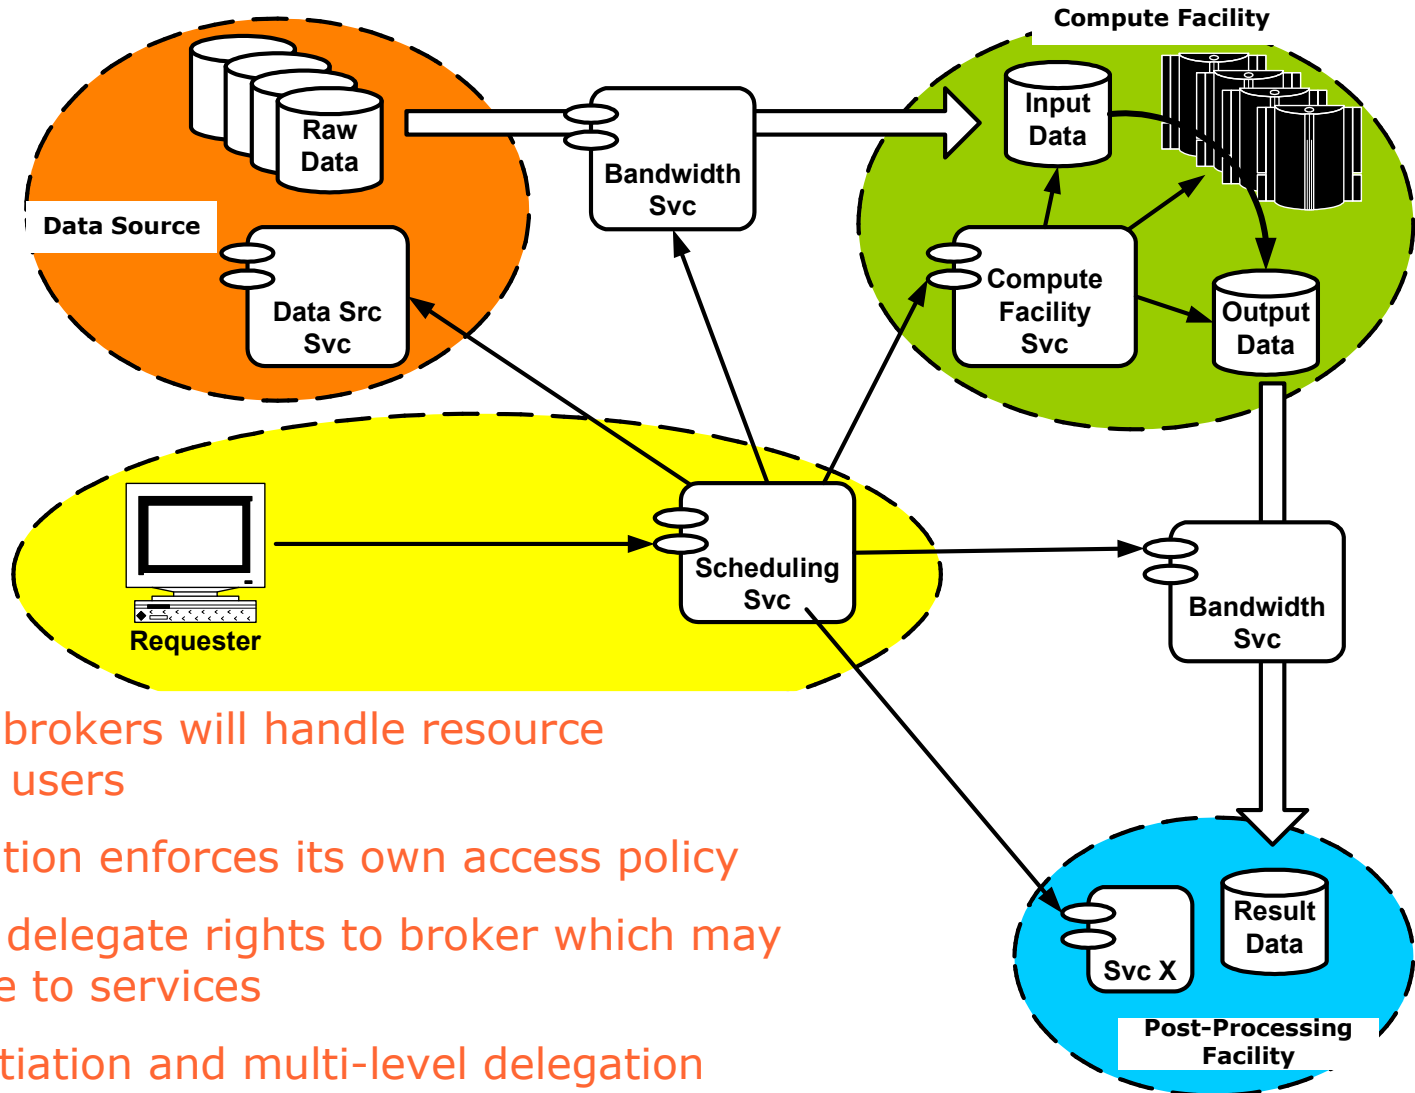
Grid Solution: Use Virtual Organization as Bridge



Virtual Organization Enables Access



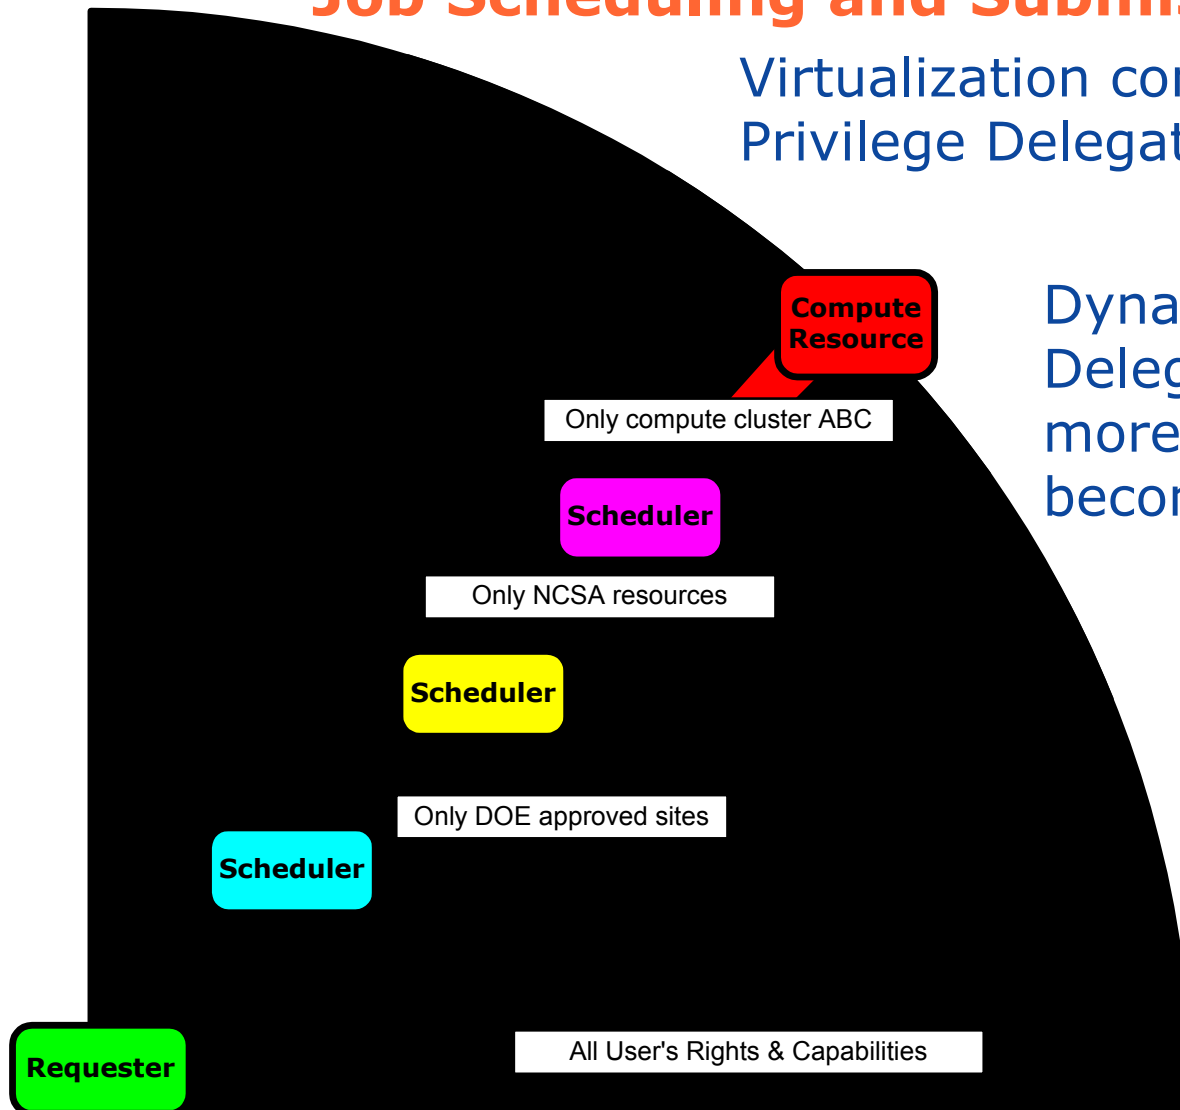
Security of Grid Brokering Services



- It is expected brokers will handle resource coordination for users
- Each Organization enforces its own access policy
- User needs to delegate rights to broker which may need to delegate to services
- QoS/QoP Negotiation and multi-level delegation

Propagation of Requester's Rights through Job Scheduling and Submission Process

Virtualization complicates Least Privilege Delegation of Rights



Dynamically limit the Delegated Rights more as Job specifics become clear

Trust parties downstream to limit rights for you... or let them come back with job specifics such that you can limit them

Dynamic Resource Management

- Compute jobs are run in newly created accounts
 - ◆ Any account creds are created on the fly...
- Dynamic account/sandbox creation
 - ◆ X.509 identity registration procedure doesn't work...
 - ◆ Identity assertion not very useful...
- Newly created key pairs are "the" identity creds
 - ◆ Only "Host" key is long-lived
 - ◆ Only "Host" can be used to derive authz from
- Currently use proxy-certs to issue authz-assertions
 - ◆ "Host" asserts that requester can be trusted by account
 - ◆ "Host" asserts account can be trusted by requester
 - ◆ Requester asserts account can work on behalf of requester

(Grid) Use Cases for Delegation of Rights

- Grid applications traverse admin boundaries
- Services work on behalf of others
- Authz-assertion chains are built dynamically
- Combination of multiple assertions decides decisions

- Need for the “right” policy language
 - ◆ with industrial strength open source toolkit
- Policy engine should be present at all control points
 - ◆ Embed engine in our Globus Toolkit


- Industry lags 2-3 years behind,
but will face the exact, same requirements...

Outline

- ProxyCertificates
 - ◆ Von Welch
- (Grid) Use Cases for Delegation of Rights
 - ◆ Frank Siebenlist
- SPKI
 - ◆ Carl Ellison
- XrML
 - ◆ Ravi Pandya
- TrustBuilder
 - ◆ Kent Seamons



TrustBuilder: Automated Trust Negotiation in Open Systems



Kent Seamons

Brigham Young University
Internet Security Research Lab
seamons@cs.byu.edu

3rd Annual PKI R&D Workshop, Gaithersburg, MD, April 12, 2004



Outline

- ◆ Trust establishment in open systems
- ◆ Overview of trust negotiation
 - Sensitive credentials and access control policies
 - Research directions
- ◆ TrustBuilder
 - TLS-based trust negotiation protocol
- ◆ Future work

Trust Negotiation Collaborators

◆ Theory

- M. Winslett, UIUC
- T. Yu, NCSU
- N. Li, Purdue
- W. Winsborough, GMU
- J. Mitchell, Stanford

◆ Systems

- K. Seamons, BYU
- C. Neuman, T. Ryutov, B. Tung, USC/ISI
- H. Orman, Purple Streak


◆ Applications

- W. Nejdl, U. Hannover
Educational consortia
- J. Basney, V. Welch, NCSA
Grid computing

◆ Funding

- DARPA (Dynamic Coalitions Program)
- NSF (ITRs on TN, disaster response)
- Industry (ZoneLabs, Dallas Semiconductor, Network Associates Laboratories)

Trust Establishment in Open Systems



- ◆ Problem: Identity is not relevant
- ◆ Solution: Access control decisions are based on attributes of both the client and server (mutual trust)
 - Client attributes: citizenship, security clearance, job classification, annual salary, affiliations, etc.
 - Server attributes: membership, privacy policy, customer satisfaction, result of recent security audit, etc.


Digital Credentials

- ◆ A credential is the vehicle for carrying attribute information reliably
- ◆ A credential contains attributes of the credential owner asserted by the issuer (attribute authority)
- ◆ Properties: verifiable and unforgeable



Credentials Sensitivity

- ◆ Credentials may contain sensitive information and should be treated as protected resources



Access Control Policies

- ◆ Credential disclosure is governed by an access control policy
 - Specifies credentials that must be received from another party prior to disclosing the sensitive credential to that party



ISRL

Internet Security Research Lab
BRIGHAM YOUNG
UNIVERSITY

Trust Negotiation

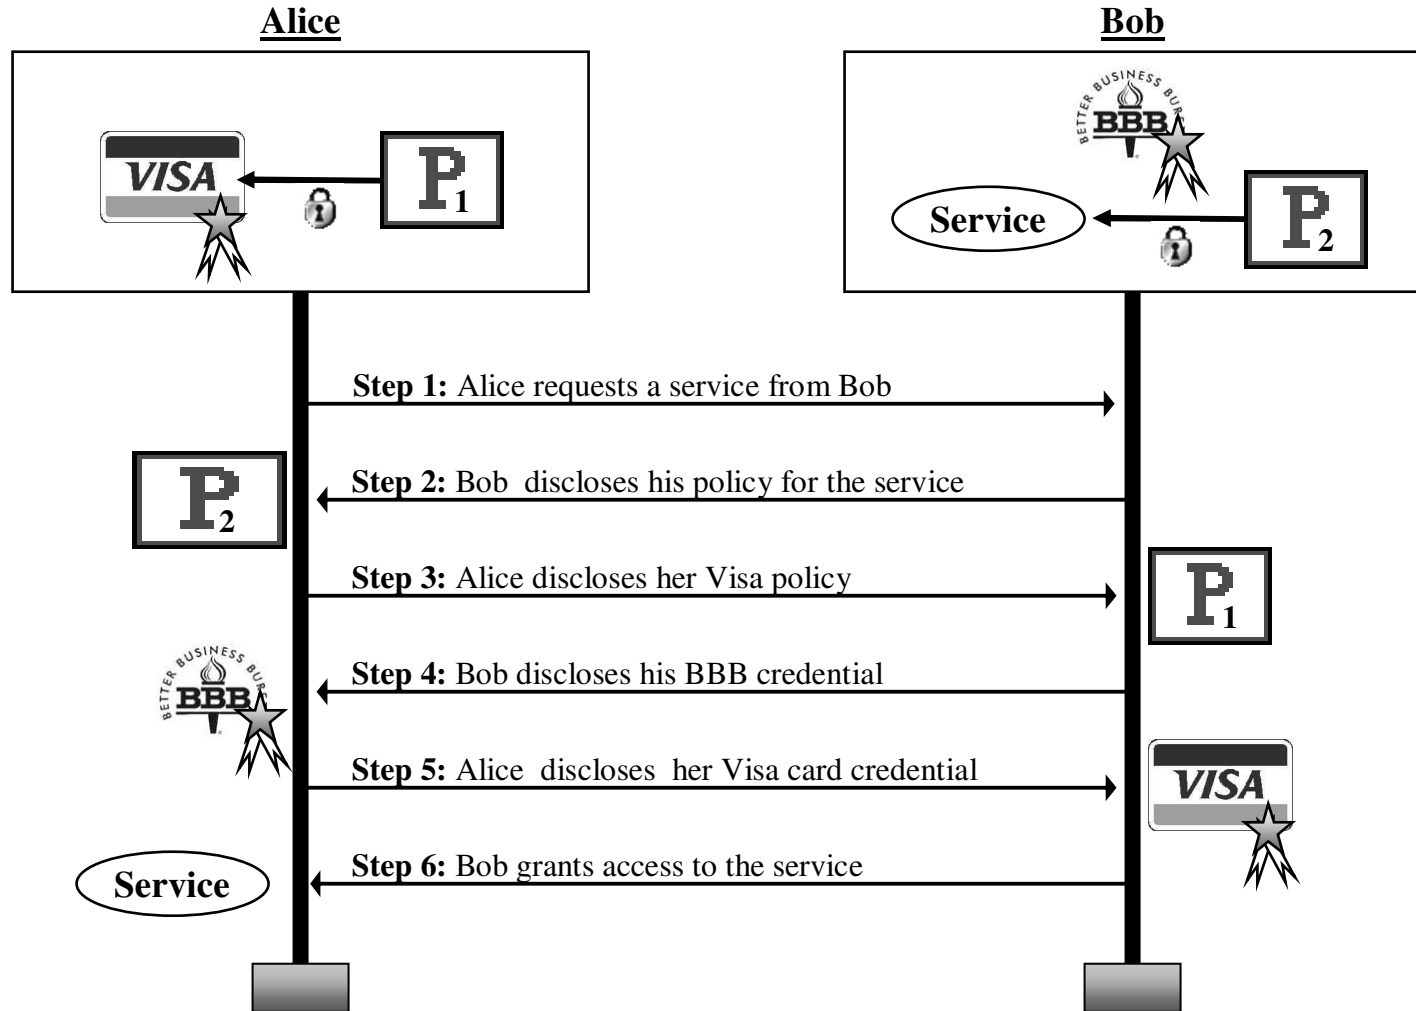
- ◆ The process of establishing trust between strangers in open systems based on the attributes of the participants



Trust Negotiation Approaches

- ◆ Naïve:
 - Disclose all credentials with each request for service
- ◆ Trial and error
 - Disclose all credentials that are not sensitive, disclose sensitive credentials after required trust is established
- ◆ Informed
 - Disclose relevant policy first, then only disclose credentials necessary for a successful trust negotiation based on the trust requirements within the policy
- ◆ Advanced cryptography
 - Demonstrate attributes without disclosing credentials

Trust Negotiation Example



Research Directions


ISRL

Internet Security Research Lab
BRIGHAM YOUNG
UNIVERSITY

- ◆ Policy languages
 - Requirements (Seamons, Winslett – Policy 2002)
 - Compliance checker requirements
 - Policy language design
 - IBM TPL – (Herzberg et al., Oakland 2001)
 - RT - (Li, Mitchell, Winsborough, Oakland 2002)
 - Delegation of attribute authority, role mappings between organizations
 - PeerTrust (Nejdl et al., ESWS 2004)
 - Policy analysis tools (Li, Winsborough, Mitchell)
- ◆ Finding credentials at run time (Winsborough, Li)
- ◆ Preventing leaks/attacks during negotiation
 - Hidden credentials (Holt et al, WPES 2003)
 - OSBE (Li et al., PODC 2003)
 - Ack policies (Winsborough et al., Policy 2002)
 - Policy filtering (Yu et al.)
- ◆ Support for sensitive access control policies (Seamons, Winslett, Yu)
- ◆ Negotiation protocols & strategies
- ◆ Wireless and mobile device architecture for trust negotiation – surrogate TN
- ◆ Testbed implementations - HTTPS, TLS, content-triggered TN, ...

TrustBuilder Architecture


- ◆ Goal – Ubiquitous trust negotiation
- ◆ TrustBuilder integrates into existing Internet technologies
 - Current Deployments - HTTPS, TLS, SSH, SMTP



Trust Negotiation in TLS (TNT)

- ◆ TLS-based protocol for trust negotiation
- ◆ Resulted from an analysis of the SSL/TLS handshake protocol for its suitability as a protocol for trust negotiation
 - TLS provides an option for client/server authentication using certificates
 - Goal: extend TLS client/server authentication to support trust negotiation

TLS Handshake Protocol using RSA Key Exchange



Limitations of TLS for Establishing Trust between Strangers

- Certificates are exchanged in plain text
- Client and server each disclose only one certificate chain
- Server can specify a list of trusted certifying authorities; client cannot
- Server always discloses its certificate first
- Server certificate ownership is not yet established when the client discloses its certificate

Extend TLS Authentication to Support Trust Negotiation

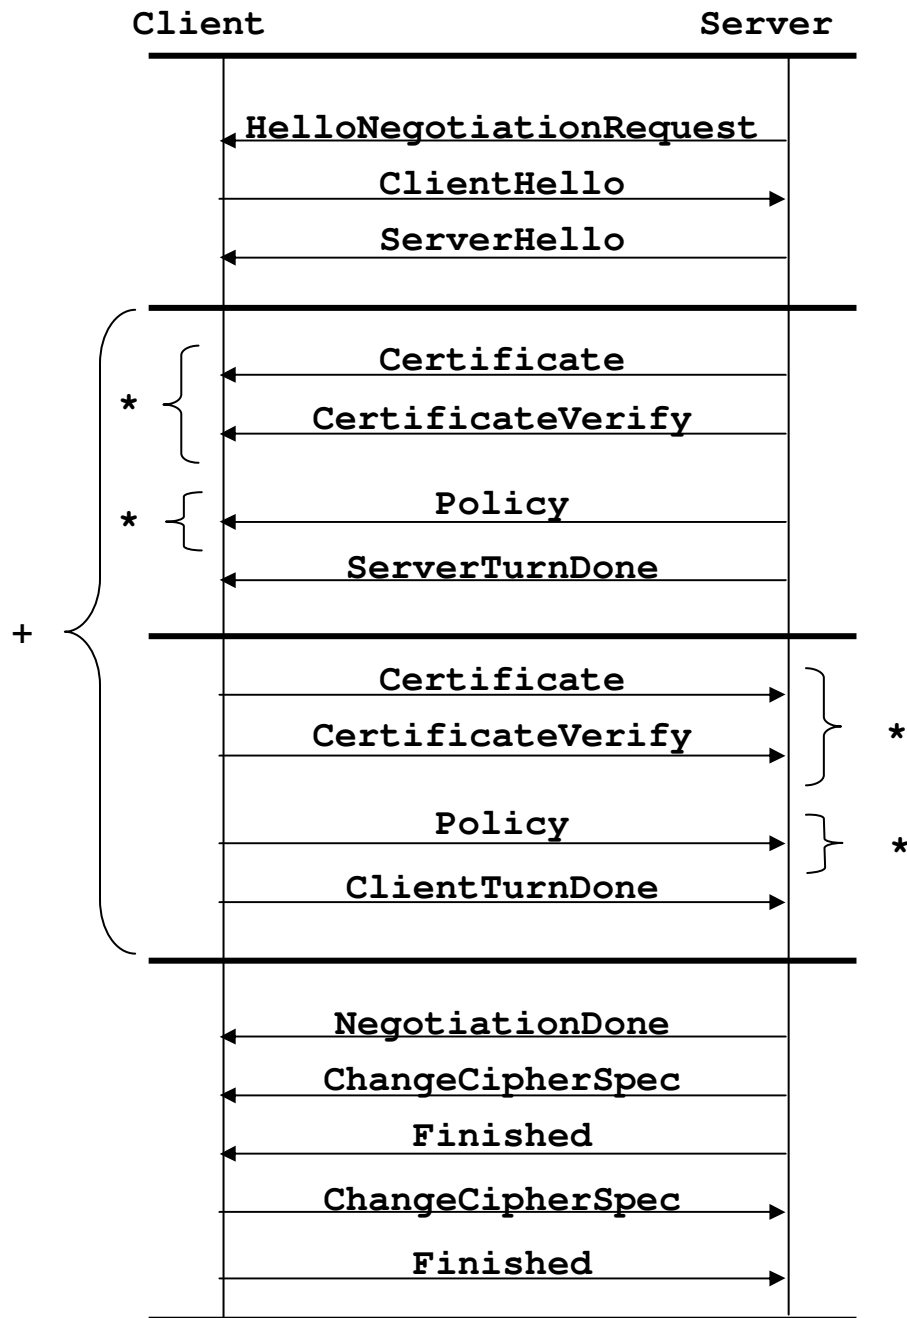
- ◆ Extend the TLS handshake protocol to function as a trust negotiation protocol
- ◆ TNT leverages existing and proposed features of the TLS handshake protocol
 - Client hello and server hello extensions
 - TLS rehandshake
 - Session resumption

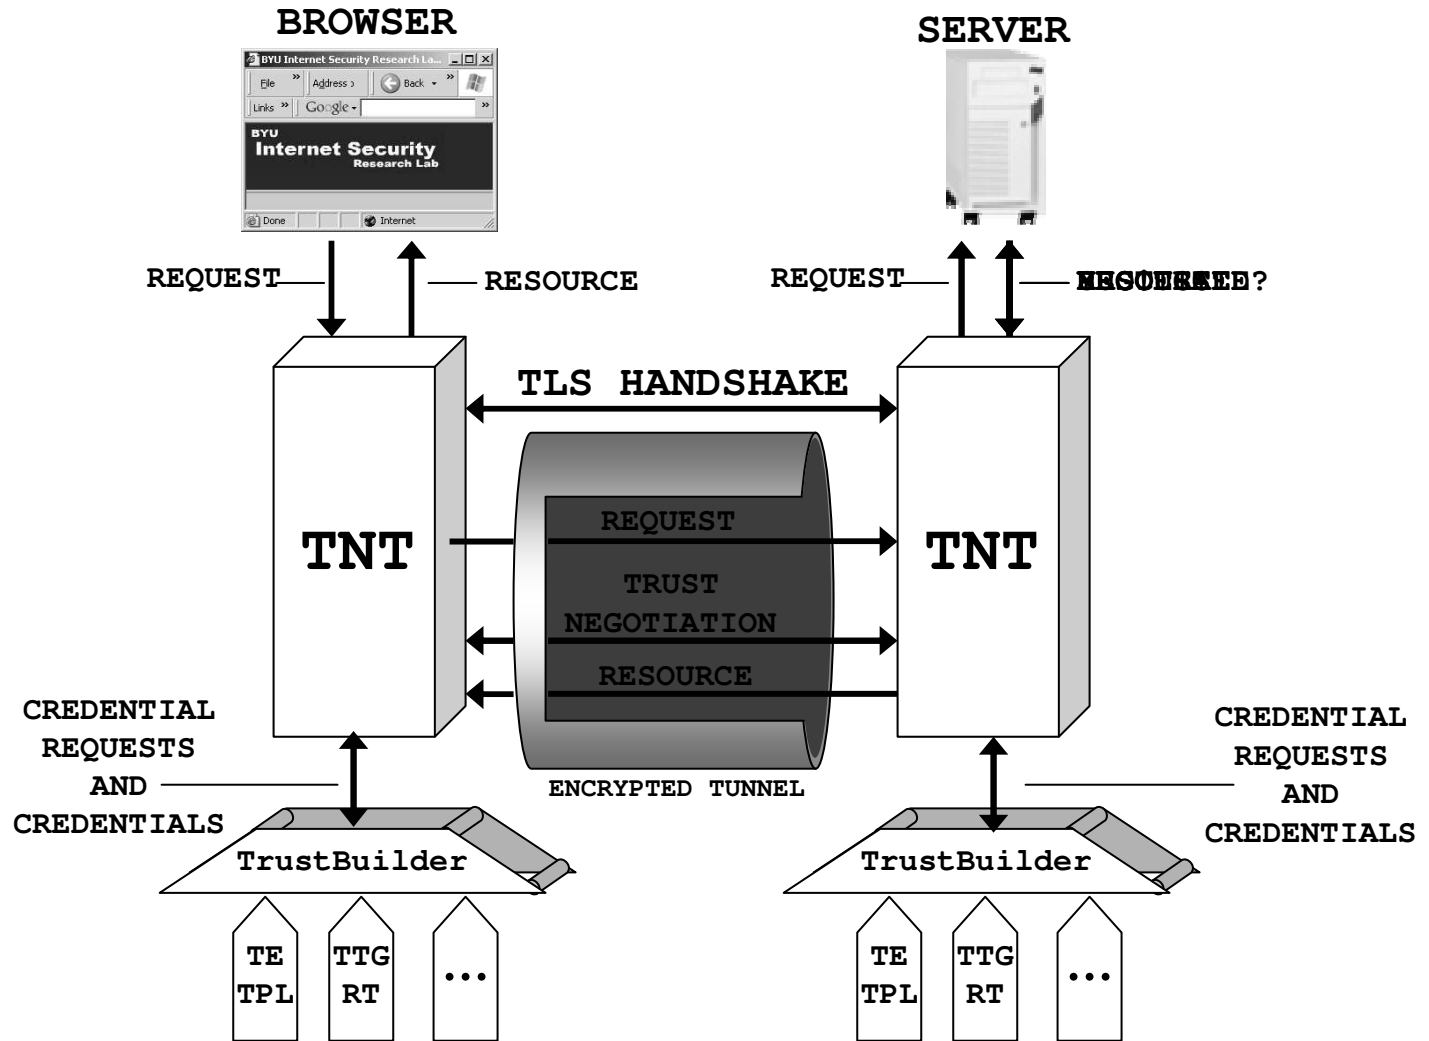
TLS Rehandshake

- ◆ In the context of an encrypted TLS session, either the client or the server may initiate a rehandshake.
 - The server desires further certificates from the client for purposes of authentication or authorization.
 - Cipher suite upgrading
 - Replenishment of keying material
- ◆ Trust negotiations involving sensitive credentials and policies must be conducted over a secure channel in order to remain confidential. The initial TLS handshake is not confidential.
- ◆ TNT is designed to occur in the context of a TLS rehandshake.

TNT Protocol

Overview:





TNT Implementation

- ◆ A prototype of TNT has been developed for the TrustBuilder architecture
 - TNT implementation is an extension to the Java PureTLS toolkit developed by Eric Rescorla (see <http://www.rftm.com/>)
 - Policy language and compliance checker is built using the IBM Trust Establishment system developed at the IBM Haifa Research Lab (RSA Security Conference 2001)

Future work

- ◆ Integrate emerging trust negotiation policy languages (RT, PeerTrust) into TNT to determine if the protocol is general purpose
- ◆ Policy creation tools
- ◆ Requirements for real-world applications
 - Grid computing, Semantic Web
- ◆ Privacy protection during trust negotiation
- ◆ Trust negotiation in Kerberos PK-INIT
- ◆ Architectures for mobile devices – surrogate trust negotiation



ISRL

Internet Security Research Lab
BRIGHAM YOUNG
UNIVERSITY

Work in Progress Session

[Ben Chinowsky](#), *Internet2*

Public Key Infrastructure (X.509) Library [libpkix]

Steve Hanna, Sun Microsystems

Steve presented libpkix (<http://libpkix.sourceforge.net>), an extensible C library for building and validating cert paths. They are looking for project participants. Various research questions are involved; one of particular importance is how you limit the amount of effort expended on pathbuilding. The Mozilla developers are very interested in this work.

The Bear Project

Sean Smith, Dartmouth College

Sean discussed Bear (<http://www.cs.dartmouth.edu/~sws/abstracts/msmw03.shtml>). This work is designed to address the question, "why should you trust computing that happens somewhere else?" For example, why should I trust a Shibboleth attribute authority to be giving my attributes only to the right people? The client wants to only have to provide a cert; the server doesn't want to spend money, and wants easy maintenance and good performance. The IBM 4758 doesn't solve the server problems, as it's expensive and awkward to code for. The Bear project attempts to provide 4758-like functionality on a standard machine equipped with version 1.1b of the TCPA/TCG TPM (e.g., many IBM NetVistas). Bear is now running with OpenCA in the lab; the code is at <http://enforcer.sourceforge.net>. Smith noted some weaknesses: Bear is probably vulnerable to power analysis; unprotected systems between the client and server will create vulnerabilities; and there are probably holes in the OS code. AEGIS could address some of these weaknesses. A revised and updated Bear paper will appear in ACSAC in December 2004.

Domain Name System Security (DNSSEC) Update

Sam Weiler, SPARTA

Finally, Sam followed up last year's WIP session on DNSSEC with a discussion of what it will take to motivate DNSSEC deployment. Security is expensive to implement, and Weiler pointed out that with security you're basically "buying brittleness" anyway. Three positions were expressed:

- Mary Thompson was the most optimistic, predicting that people will adopt DNSSEC once the technology is mature.
- Steve Hanna suggested spam as a driver, using DNSSEC to authenticate senders or MTAs. Also, some communities (e.g. ISPs) might be interested in authenticating messages to know that they're coming from within the community.
- John Linn suggested that as security is largely about assurance that information is really accurate, well-publicized DNS spoofing might get people appropriately worried. This position found the greatest resonance in the group as a whole. No one knows of any good data on how much DNS hijacking there is; it was observed that security on the web has been driven forward because there have been attacks on web sites that have cost people money, and these attacks have been well-publicized. With DNS, the attacks have probably happened, but they have not been publicized. Neal McBurnett noted that one tactic that's been used to raise security awareness is to listen to the network at a conference and publicize all the passwords discovered; maybe we need to do something similar with DNS, using a tool such as dnsspoof.

How to build a PKI that works

Peter Gutmann
University of Auckland

How to build an X.509 PKI that works

Peter Gutmann
University of Auckland

Preliminaries

Whose PKI are we talking about here?

- Not SSL certs
 - Certificate manufacturing, not PKI
 - It's just an expensive way of doing authenticated DNS lookups with a TTL of one year. Plenty of PK, precious little I — Peter Gutmann on the crypto list
- Not PGP, SPKI, *ML, etc
 - Doing fairly well in their (low-I) area
- Not government PKI initiatives
 - Government IT project reality distortion field, keep pumping in money until it cries Uncle
 - Even then, the reality distortion has failed in parts of Europe, Australia

Preliminaries (ctd)

This is PKI for the rest of us

- Businesses, individuals, etc

Talk covers exclusively technical issues

- Policies are someone else's problem

Ted says that whenever he gets asked a religious question he doesn't understand he always responds with "Ah, that must be an ecumenical matter" which universally produces nods of admiration at the profound wisdom of the statement. It seems that that the PKIX list equivalent is "Ah, that must be a policy matter"

— Father Ted (via Anon)
- Some religion may sneak in

Preliminaries (ctd)

Microsoft bashing: An apology in advance

- Their PKI software is the most widespread, and features prominently in examples because of this
- There is no indication that other software is any better, it just gets less publicity

It may be a little controversial...

56th IETF agenda item, submitted as a joke when someone pointed out that PKIX didn't have any agenda

What needs to be done to make PKI work?

This forum will be open to all PKIX members, and will constitute a large pool filled knee-deep with custard.

Marquis of Queensberry Rules, but with pies substituted for gloves. Participants are expected to provide appropriate clothing. Remaining IETF members will look on in amusement or dismay, depending on their views on PKI

Meeting minutes at

<http://www.cs.auckland.ac.nz/~pgut001/misc/minutes.txt>

Why do we need “a PKI that works”?

PKI is in trouble

PKI is ‘Not Working’ (Government Computing, UK)

“Trust and authentication has been a huge problem for us. We haven’t got a solution for authentication. We’ve been trying with PKI for about 10 years now and its not working because it’s a pain to implement and to use”.

Billion Dollar Boondoggle (InfoSecurity Mag, US)

A recent General Accounting Office report says the federal government’s \$1 billion PKI investment isn’t paying off. [...] The GAO says widespread adoption is hindered by ill-defined or nonexistent technical standards and poor interoperability [...] Despite stagnant participation, federal officials are continuing to promote the [PKI].

PKI is in trouble (ctd)

Gatekeeper goes Missing (The Australian)

Five years after then finance minister John Fahey launched Gatekeeper to drive public and business confidence in e-commerce, government department and agency interest in PKI is almost zero.

A spokesperson for the Attorney-General's Department said: "I am very grateful for the fact that none of my colleagues has come up with a good use for it. When they do, I will have to do something about it".

End of the line for Ireland's dotcom Star (Reuters)

The company would have done better to concentrate on making its core PKI technology easier to deploy, a shortcoming that became a key reason Baltimore's UniCERT PKI technology never went mainstream.

PKI is in trouble (ctd)

International and New Zealand PKI experiences across government (NZ State Services Commission)

Based upon overseas [Australia, Finland, Germany, Hong Kong, US] and New Zealand experiences, it is obvious that a PKI implementation project must be approached with caution. Implementers should ensure their risk analysis truly shows PKI is the most appropriate security mechanism and wherever possible consider alternative methods.

PKI's Image Problem

The message to potential users from mainstream media coverage: PKI doesn't work

...as computer security professionals, we feel that it is our duty to advise the legislature of the critical importance of requiring the use of a PKI for this system, preferably with multiple root CAs and online certificate revocation.

— Cryptographer John Kelsey proposing a means of killing a DRM initiative by the Copyright Policy Branch of Canadian Heritage

Why is PKI in trouble?

The usual suspects...

- Difficult to deploy
- Expensive
- Hard to use
- Lack of interoperability
- Poor match to pressing real-world problems
- Etc etc etc

The PKI Grand Challenge

Get the basic infrastructure in place before we worry about chrome tailfins, fuzzy dice, certificate warranty permanent qualifier policy logotype extensions, ...

- I can add theme music to my certificate if I want, but the only way to publish it is to stick it on my home page
- There'll be plenty of time to add the fuzzy dice once the basic infrastructure is in place

I think a lot of purists would rather have PKI be useless to anyone in any practical terms than to have it made simple enough to use, but potentially "flawed"

— Chris Zimman

I still can't use PKI to authenticate myself for the PKI Workshop...

PKI Grand Challenges

Challenge #1: Key lookup

- Original PKI was Diffie and Hellman's "Public File" in 1976
- In 1976, I couldn't look up your public key online
- After thirty years' work, I still can't look up your public key online

Challenge #2: Enrolment

- A torture test for users to see how badly they really want a cert
- Pain of enrolment leads to terrible key hygiene

Challenge #3: Validity checking

- Real-time check to match expectations of online banking, share-trading, bill payment, etc etc

PKI Grand Challenges (ctd)

Challenge #4: User identification

- X.500 DNs (enough said)
- Mostly solved in a de facto manner

Challenge #5: No quality control

- You cannot build a product so broken that it can't claim to be X.509
- Users *notice* that things don't work → PKI image problem (see challenge #6)

PKI Grand Challenges (ctd)

Challenge #6: Implementor / user apathy (HCI)

- Complexity / lack of understanding ↔ lack of motivation to do things right
 - Example: Re-checking certificate against an old CRL on disk meets requirements for a revocation check
- Current designs make it too easy to just go through the motions

Well, that's a nice theory, but...

It's practice, not theory

- Based on extensive user feedback / usability testing
- Refined over many years
- Designed to maximise ease of use, correct functionality
 - You have to really work hard to get it wrong
- Designed to minimise implementer pain

This is not just a gedanken experiment / unproven hypothesis

Challenge #1

Key Lookup

Pre-history of Key Lookup (and Certs)

Original 1976 paper on public-key encryption proposed the Public File

- Public-key white pages
- Key present → key valid
- Communications with users were protected by a signature from the Public File

Not very practical in 1976

- Key lookup over X.25?
 - Having to interrupt a circuit-switched connection to do a Public File lookup was the original motivation for offline certificates (1978)
- A very sensible, straightforward approach now that there's a WWW

The Key Lookup Problem

The problem

- Get me `joe@foo.com`'s key(s)
- Get me `foo.com`'s key(s)

Clayton's solutions: S/MIME, SSL

- Send out all your certificates with each message
- Lazy-update distributed key management

The Web as the Public File

We have a Public File

- It's called the WWW
We have a system, it is called the Web, everyone else lost, get over it
— Phillip Hallam-Baker

Quick-n-dirty solution: Google

- Stick a base64-encoded certificate on your home page
- Add a standardised string for search engines,
`certificate joe@foo.com`
- Google, cut & paste
- Clunky, but simple and effective
 - Better than anything else we have today

The Web as the Public File (ctd)

Proper solution: Use HTTP to fetch keys

- `GET uri?attrib=value`
`GET /search-cgi?email=joe@foo.com`

ID types required

- S/MIME, SSL/TLS, IPsec, PGP, SIP, etc
 - Email, domain name, URI
- Cert chaining
 - Issuer DN, keyID
- S/MIME
 - issuerAndSerialNumber
- PGP
 - PGP keyID

Implementation

HTTP glue + anything you want

- Berkeley DB
 - Lightweight { key, value } lookup
- RDBMS
 - ODBC is built into every copy of Windows
 - ODBC glue for most Unix systems
 - MySQL or Postgres is built into most copies of Linux
 - JDBC for Java
 - Ties into existing corporate databases (SQL Server, Oracle)
- ISAM
- Flat files
 - c.f. PGP's HKP servers
- X.500 / LDAP if you insist

Implementation (ctd)

Implementation effort

- MySQL (server): 30 minutes
 - Every database on the planet is already web-enabled
 - This is what many web servers do all day long
- Java (server): A few hours
- Visual Basic (client): About 5 minutes

Lightweight client

- ~100 lines of code on top of TCP/IP stack in an embedded network device

Other Features

Pre-construct URLs for certificates

- Print on business cards
- Help-desk can mail to users who can't find their certificates
- Enforce privacy by perturbing the search key
x-encryptedSearchKey=...
- Enforce access controls by authenticating the search key
x-macSearchKey=...

Other Features (ctd)

Standard techniques used to manage high loads

- It's a standard web server with static pages
 - Web101
- If Amazon / CNN.com can handle this...

More details / rationale in “Certificate Store Access via HTTP”

But what about X.500 / LDAP?

If you can't be a good example then at least you can be a horrible warning

But what about X.500 / LDAP?

So far, LDAP has not done a great job of supporting PKI requirements.

— Steve Kent, PKIX WG chair

The X.500 linkage [...] has led to more failed PKI deployments in my experience than any other. For PKI deployment to succeed you have to take X.500 and LDAP deployment out of the critical path.

— Phillip Hallam-Baker, Verisign principal scientist

- If you really want to, you can always use X.500 / LDAP as another backend for the HTTP certstore — it's not picky

The most effective way I've found to search an X.500 directory to locate a certificate is by Internet email address

— PKI developer

Challenge #2

Enrolment

What it should be like: The DHCP Model

User wants to use TCP/IP / email / WWW

- DHCP client automatically discovers the server
- Client requests all necessary information from the server
- Auto-configures itself using returned information
- User is online without even knowing that the DHCP exchange happened

What it is like: The X.25 Model

User is required to use X.25

- Dozens of parameters to manually configure
- Different vendors use different terms for the same thing
- Get one parameter wrong and nothing works
- Problem diagnosis: Find an X.25 expert and ask for help

The vast majority of users detest anything they must configure and tweak. Any really mass-appeal tool must allow an essentially transparent functionality as default behaviour; anything else will necessarily have limited adoption.

— Bo Leuf, *Peer to Peer*

How bad is it really?

Obtaining a certificate from a large public CA

- User had to ask where to get the certificate
- Filled out eight (!!) browser pages of information
- Several retries due to values being rejected, had to ask for help several times, searched for documentation such as a passport, etc etc
- Cut & pasted data from emailed message to web page
 - Multiple random strings had to be manually copied over
 - Emailed cookies: Only one should be necessary

How bad is it really? (ctd)

- Filled out more fields in eleven further web pages
 - Much of the contents were incomprehensible to the user:
“certificate Distinguished Name”, “X.509 SubjectAltName”
My grandmother just won’t understand the meaning of
“initial-policy-mapping-inhibit” no matter how much she
loves me.
 - David Cross on ietf-pkix
 - User guessed and clicked “Next”
- Web page announced that a certificate had been issued, but none seemed available

How bad is it really? (ctd)

- Emailed message provided a link to click on
- More web pages to fill out
- Switch to another browser to download file
- Clicking on the file had no effect

At this point the user gave up

How bad is it really? (ctd)

Time taken: > 1 hour (with outside assistance)

- Usenet posts/email suggest that most skilled technical users take between 30 minutes and 4 hours to get a certificate
“There’s a myth [...] that the issuance of a public certificate is a remarkably heavyweight operation. You know, you must need steam-powered equipment in the basement of your facility in order to stamp out those certificates, which have to be made out of titanium or what have you”

— Matt Blaze, Security Protocols Workshop

The Machine that Issues Certificates,

[http://www.cs.auckland.ac.nz/
~pgut001/misc/certificates.txt](http://www.cs.auckland.ac.nz/~pgut001/misc/certificates.txt)

Consequences of enrolment difficulties

Pain of enrolment encourages poor key hygiene

- Company spends \$495 and several hours’ work creating a key and getting a Verisign certificate for it
- Most practical (in terms of time and money) application of this is to re-use it everywhere
 - “It cost us \$xxx/yyy hours’ effort to get this key, we’re not going through all that again”

Much of the problem is social/financial

- Certificates are expensive to obtain
- Certificates are troublesome to obtain
- Users are given a considerable incentive to re-use certs/keys

Consequences of enrolment difficulties (ctd)

CAs generate private keys for users and mail them out as PKCS #12 files

- Password is sent as separate mail or is easily guessed (8 characters, uppercase-only)
- This is standard practice for many, many CAs

I didn't generate PKCS #10. My CA does not support this request [...] CA sends me two files – private key and certificate.

the certificates and the key pairs are centrally generated and send to the user as PKCS#12 files. The user imports this file in his Internet Explorer and can use it for SSL client authentication. This works successfully.

continues...

Consequences of enrolment difficulties (ctd)

CA generates only PKCS12 key files [...] I can not find an exact explanation how to read a PKCS12 private key form such a file.

Plus, they attach your certificate AND _private key_ to the bottom of the message. The idea is that you copy and paste the cert + private key into a file for the client API to use when it connects. Basically, they are sending all of the information [...] through plain, unencrypted, email.

I have two files from CA – private key and certificate.

what is the format to use for sending me a private key certificate when the CA does the whole process themselves - and want to send me a pin code and a PKCS#12 cert

continues...

Consequences of enrolment difficulties (ctd)

The CA generates an encryption key pair for the client and issues a certificate for the public key. The CA sends the private key.

import pkcs#12 files (including private key) onto the smartcard [...] Sometimes they let you even generate keypair(s) on the card and have the public part certified by the CA's, which is not always a good idea...

— Representative sampling from newsgroups and mailing lists

- One development group took to referring to the private key as “the lesser-known public key”

Consequences of enrolment difficulties (ctd)

CAs distribute their own private keys as PKCS #12 files

- The theory is that once installed, it makes the CA key trusted
- This “solution” is so common that it's warned about in the OpenSSL FAQ
- At least one computer security book contains step-by-step instructions on how to distribute your CA's private key to all users

Application developers send PKI software developers their private keys during debugging

- Verisign Authenticode code-signing keys, banking keys, etc etc

Consequences of enrolment difficulties (ctd)

Smart cards store private keys internally and don't reveal them

- “How can I use a smart card if I can't get at the key?”
what is the point in jailing the private key for life in a single smart card? This argument is totally contrary to logical thinking.
— Anon on ietf-pkix

Consequences of enrolment difficulties (ctd)

- Attempted fixes are to...
 - Construct mechanisms for sharing cards across multiple machines
 - Generate the key externally and keep a copy after it's loaded onto the card
 - Exacerbated by the mail-a-PKCS12 approach to certification
- Maybe the inconvenient fact that they keep private keys private is why crypto smart cards aren't taking off

What should enrolment be like?

The mom test: Could your mother use this?

The ISP model

- Call ISP with credit card
- ISP provides username and password
- Enter username and password, click OK
- DHCP does the rest

PKI enrolment should be similar

- Others have debugged the process for us
- Users have been conditioned to do this
- Most users can handle this

Assumptions

Basic networking services are present

- The user has a net connection, IP address, etc etc (DHCP at work)

Assumptions (ctd)

The user has some existing relationship with the certificate-issuing authority

- Issuing identity certificates to strangers doesn't make much sense
- Online banking / tax filing / loyalty program sign-up is usually handled by
 - In-person communications
 - (Snail) mailed authenticator
 - Phone authorisation
- Follows existing practice
 - People are used to it
 - Established legal precedent

Assumptions (ctd)

We're not designing a system to handle nuclear weapons launch codes

- The system need only be as secure as the equivalent non-PKI alternative
 - Techies tend to go overboard when designing authentication systems
 - Operations where a cert might be used (online banking, shopping, tax filing) all get by with a username and password
 - If it's good enough when used without certificates, it's equally good with them
- Cumbersome technology will be deployed and operated incorrectly and insecurely, or perhaps not at all
- Ravi Sandhu, *IEEE Internet Computing*

PKI Service Location

DHCP

- Limited to local subnet
- Would require modifying all existing DHCP servers
- Unnecessarily low-level: Higher-level network infrastructure is already in place

DNS SRV

- Easily added to existing servers
 - Single line in a config file
- Not supported in Win'95/98/ME
- Those who need it most don't have it
 - Expecting Auntie Ethel to install `bind` is probably a bit much

PKI Service Location (ctd)

SLP

- Service Location Protocol, specialised service-location mechanism
- Rarely used, requires configuring and maintaining yet another server/service

UPnP

- Very complex
- Requires XML (SOAP), HTML GUI interface, etc etc
- Many sites block UPnP for the same reason that they block NetBIOS

PKI Service Location (ctd)

Jini

- Very complex
- Tied to Java-specific mechanisms (RMI, code downloading, etc etc)

Others: Salutation, Rendezvous, ...

- See SLP

PKI Service Location (ctd)

Faking it

- Use of “well-known” locations for services
- Full IP service (e.g. PC): Use “pkiboot” at start of domain name
 - `foo.domain.com` → `pkiboot.domain.com`
 - Example: Corporate/organisational CA certifying users
- Partial IP service (e.g. web-enabled embedded device): Append “pkiboot” to device’s IP address or location:
 - `192.0.0.1` → `http://192.0.0.1/pkiboot/`
 - Example: Print server certifying printers
- Use HTTP redirects if necessary
- Somewhat clunky, but can be done automatically/transparently

PKIBoot: Obtaining Initial Certificates

Establishing the initial trusted certificate set (PKI TCB)

- Browsers contain over 100 hardcoded certificates
 - Unknown CAs
 - Moribund web sites
 - 512-bit keys
 - No-liability certificates
 - Keys on-sold to third parties
- Any one of these CAs can usurp any other CA
 - Implicit universal cross-certification
 - Certificate from “Verisign Class 1 Public Primary Certification Authority” could be issued by “Honest AI’s Used Cars and Certificates”
 - Browser trusts Verisign and Honest AI equally

PKIBoot: Obtaining Initial Certificates (ctd)

Why do browsers do this?

- Prime directive: Don’t expose the users to scary warning dialogs
- One-size-fits-all browser can’t know in advance which entities the user has a trust relationship with
 - Need to include as many certificates as possible to minimise the chances of users getting scary warning dialogs
 - The ideal user-friendly situation would be to automatically trust all certificates

Goal: User should only have to trust certificates of relevance to them (minimised TCB)

PKIBoot: Obtaining Initial Certificates

Initial state: No certificates

Use username + password to authenticate download of known-good/trusted certs (PKIBoot)

- Messages are protected using a cryptographic message authentication code (MAC) derived from the password
- User → PKI service: Send known-good certificates
 - User request is authenticated with MAC
- PKI service → user: Known-good certificates
 - PKI service response is authenticated with MAC
- Since only the legitimate service can generate the MAC, certificate spoofing isn't possible

Obtaining User Certificates

Initial state: CA certificates

Use MAC to authenticate the request for a signing certificate

- User → PKI service: Sign this for me
 - User request is authenticated with a MAC
- PKI service → user: Signed certificate
 - PKI service response is authenticated with a signature from the PKIBoot cert

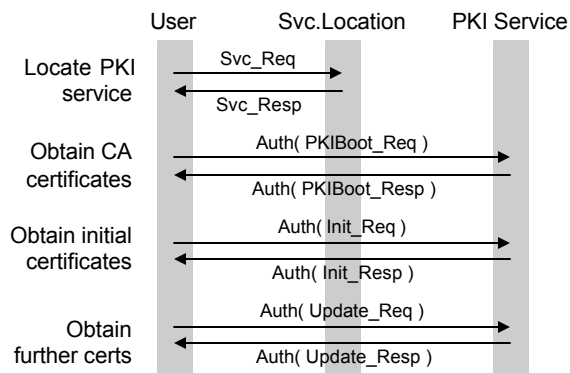
Obtaining User Certificates (ctd)

Initial state: CA certificates, signing certificate

Use signing certificate to authenticate the request for an encryption certificate

- User → PKI service: Sign this for me
 - User request is authenticated with the signing cert
- PKI service → user: Signed certificate
 - PKI service response authenticated with a signature from the PKIBoot cert

Sequence of Operations



Multi-phase bootstrap

- MAC → CA cert, signing cert request
- CA cert → response
- Signing cert → encryption cert

PnP PKI in action

User

- Enters username + password (identifier + authenticator)
 - No need to even mention certificates

Software developer

- Creates PnP PKI session
- Adds file/smart card for key storage
 - Card can be pre-personalised with enrolment information
- Adds username + password

PnP PKI in action (ctd)

PnP PKI session

- Performs PKIBoot using username + password
- Generates signing key
- Requests signing certificate using username + password
- Generates encryption key
- Requests encryption certificate using signing certificate
- Updates file/smart card with signing, encryption keys and user and CA certificates

User/Software developer

- Has keys and certificates ready for use

HCI Aspects of PnP PKI

Minimalist enrolment (with pre-personalised smart card)

- Insert card
- Enter PIN to unlock/access card
- Wait a few seconds
- Done

Enforces best practices by default

- Minimal set of trusted certificates (TCB)
- Locally-generated private keys
 - Keys can be generated inside crypto hardware
- Distinct encryption and signing keys

Details / rationale in “Plug-and-play PKI: A PKI your mother can use”

Challenge #3

Validity Checking

Current Approaches

Ignore it entirely

Go through the motions

- Repeatedly re-check a day / week-old CRL

OCSP

- If fed a freshly-issued cert, can't say "It's valid"
 - If fed an Excel spreadsheet (or a forged cert), can't say "It's not valid"
 - No scalability
 - Vendors eliminate replay-attack protection in order to get usable performance
- The changes we are making to scale our OCSP responder will result in the discontinuation of the nonce extension
- Verisign

What's Needed

The web has conditioned users to expect live, real-time status updates

- ebay bidding
- Amazon.com et al
- Stock trading
- Online bill payment
- Travel booking
- Paypal

Certificate validation checking should be no less functional than these systems

What's Needed (ctd)

The target: Yes/no response in as close to real-time as possible

Learning in 80 ms that the cert was good as of a week ago and to not hope for fresher information for another week seems of limited, if any, utility to us or our customers.

— PKI architect

Implementation

Query: hash(cert)

- Cert fingerprint / thumbprint recognised by any PKI software

Response: CMS(yes | no)

- Signed response (slow)
- MAC'd response (fast)
- Plain response (over secure link, very fast)

Totally unambiguous response, in real time

- It's valid right now
- It's not valid right now
 - Can be embellished with reasons, dates, etc etc

Performance

A single PC can saturate a 100Mbps link

- Connectionless (UDP) queries
 - Both queries and responses are tiny
- $O(1)$ hash table / CAM lookup
 - Query is pre-hashed by the client
- memcpy result data
 - ASN.1, but fixed format, requires no en/decoding
- Drop MAC or sig. into fixed location

You cannot build a faster validity checking mechanism

Performance (ctd)

Performance options

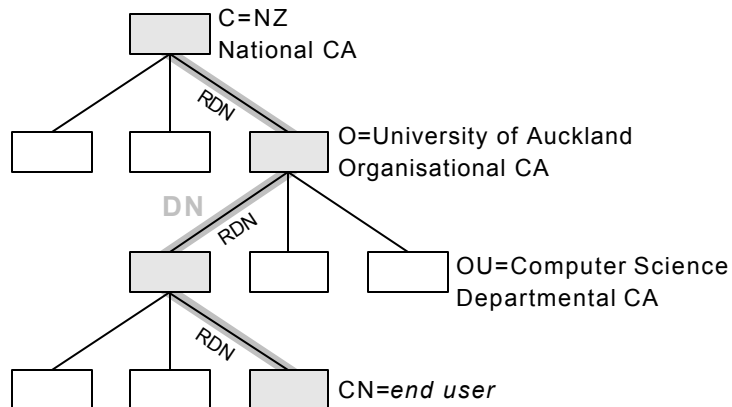
- Software-only, MAC'd response
 - Can saturate 100Mbps link
 - CMS can bootstrap MAC keys via PKC exchange
 - Key exchange can be initiated by the server to reduce load
- Broadcom 582x, scatter/gather operation
 - 4K signed responses/sec (10Mbps)
- Cavium Nitrox, all ops done on-chip
 - 40K signed responses/sec, (100Mbps)

Challenge #4

User Identification

The X.500 DN

X.500 introduced the Distinguished Name (DN), a guaranteed unique name for everything on earth



X.500 Naming

Typical DN components

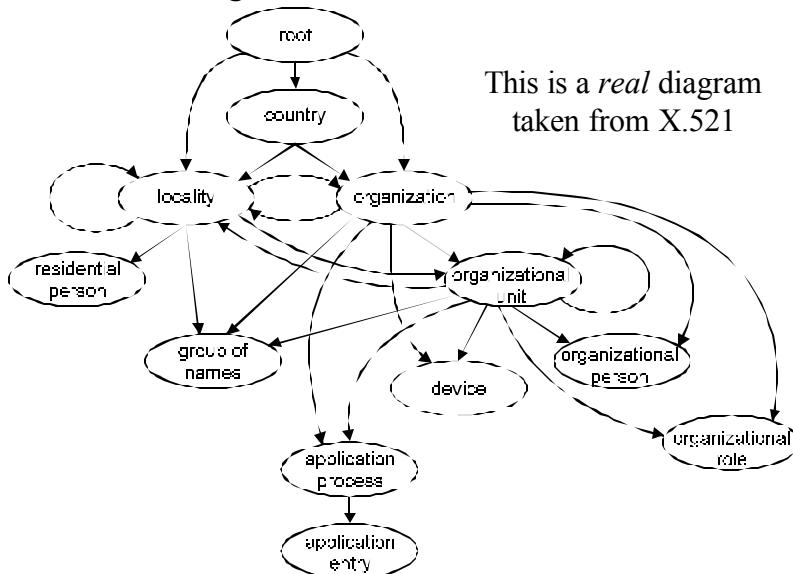
- Country C
- State or province SP
- Locality L
- Organisation O
- Organisational unit OU
- Common name CN

When the X.500 revolution comes, your name will be lined up against the wall and shot

C=US, L=Area 51, O=Hangar 18, OU=X.500 Standards Designers, CN=John Doe

Problems with X.500 Names

No-one ever figured out how to make DNs work



Problems with X.500 Names (ctd)

No clear plan on how to organise the hierarchy

- Attempts were made to define naming schemes, but nothing really worked
 - NADF
- People couldn't even agree on what things like 'localities' were

Hierarchical naming model fits the military and governments, but doesn't work for businesses or individuals

Problems with X.500 Names (ctd)

DNs provide the illusion of order while preserving everyone's God-given Freedom to Build a Muddle

Sample problem cases

- Communal living (jails, boarding schools)
- Nomadic peoples
- Merchant ships
- Quasi-permanent non-continental structures (oil towers)
- US APO addresses
- LA phone directory contains > 1,000 people called "Smith" in a nonexistent 90000 area code
 - A bogus address is cheaper than an unlisted number
 - Same thing will happen on a much larger scale if people are forced to provide information (c.f. cypherpunks login)

Problems with X.500 Names (ctd)

For a corporation, is C, SP, L

- Location of company?
- Location of parent company?
- Location of field office?
- Location of incorporation?

For a person, is C, SP, L

- Place of birth?
- Place of residence/domicile?
 - Dual citizenship
 - US military personnel can choose “resident” state for tax purposes
 - Stateless persons
 - Nomads
- Place of work?

DNs in Practice

Public CAs typically set

C = CA country or something creative (“Internet”)

O = CA name

OU = Certificate type / class / legal disclaimer

CN = User name or URI

email = User email address

- Some European CAs add oddball components required by local signature laws
 - Italy adds IDs like BNFGRB46R69A944C

DNs in Practice (ctd)

- Some CAs modify the DN with a nonce to try and guarantee uniqueness
 - Armed services CA adds last 4 digits of SSN
 - Another CA encodes random CA/RA-specific data
The disambiguating factor will be variable length alphanumeric [...] for example: XYZ221234 [...] or, for example ABC00087654321.
 - GTE Government Systems Federal PKI pilot

Some DNs are deliberately mangled

For educational institutions here in the US, FERPA regulations apply. The way we do this here at Wisconsin is to only include a bunch of random gibberish in the DN as an identifier.

— Eric Norman on ietf-pkix

DNs in Practice (ctd)

Private CAs (organisations or people signing their own certificates) typically set any DN fields supported by their software to whatever makes sense for them

- Some software requires that all of { C, O, OU, SP, L, CN } be set
 - “Invent random values to fill these boxes in order to continue”
- Resulting certificates contain strange or meaningless entries as people try and guess values, or use dummy values
 - “... a bunch of random gibberish in the DN...”

DNs in Practice (ctd)

The goal of a cert is to identify the holder of the corresponding private key, in a fashion meaningful to relying parties.

— Steve Kent

- Minimalist DNs
 - “Fred’s Certificate”
 - “My key”
 - “202.125.47.110”

DN Encodings

Encoding of DN is more or less random

- Arbitrary grouping of AVAs, ordering and number of RDNs, etc etc

DNs may be encoded backwards

- A side-effect of the RFC 1779 string representation
- Java-created certs often have backwards DN because of this
- Some .NET DN orders are forwards, some backwards
 - GetIssuerName / GetSerialNumber vs. MMC snap-in
- One European national CA encodes DN backwards and forwards at random
 - Other CAs are more consistent in getting DN backwards

DN Encodings (ctd)

Applications enforce arbitrary limits on data elements
(GCHQ/CESG interop testing)

- Number/size of DN elements
- Size of encoded DN
- Ordering/non-ordering of DN elements
 - Allow only one attribute type (e.g. OU) per DN
 - Assume CN is always encoded last

The real DN encoding / name comp.rules

There is no name comparison rule but binary compare, and `memcmp()` is its implementation

- Originator encodes the DN any way they want
- Further “re-encoding” is done via `memcpy`
- Comparisons are done via `memcmp`

While technically there's this DN compare algorithm in RFC2459 or the evil X.500 version anyone with any sense ignores it completely and treats DNs as equal only if they have the same encoding.

— PKI developer

We treat DNs as opaque blocks of binary data [...] we yank the exact binary blob out of the certificate and combine that with the exact binary blob of the serial number.

— S/MIME developer

The real DN encoding / name comp.rules (ctd)

These are the only rules that always work

- No matter how garbled the DN, they'll handle it
- Performing a bit-for-bit copy ensures that other apps get to see exactly what they need to see

We are testing signing and encryption in S/MIME software [...]

It seems that all the software we have tested (eg. MSoft, Utimaco) tend to do somekind of binary comparison on the certificate.

— Saku Vainikainen on ietf-pkix

The real DN encoding / name comp.rules (ctd)

Application developers learn these rules fairly quickly

- Client submits cert request with PrintableString
- CA returns cert with UTF-8 String
- Client app rejects the cert because the DN doesn't match

“Don't user Master Documents in MS Word”

“Don't change the monitor frequency settings in XF86Config”

“Don't rewrite DN's in certificates”

— Peter Gutmann on ietf-pkix

Challenge #5

Quality Control

Quality Control: The absence thereof

You can't build an app so broken that it can't claim to be X.509

- Any old rubbish can claim to be X.509, and frequently does

The X.509 brand has been diluted to the point of worthlessness

- (Deeply-buried) PGP has been sold as X.509
- “The other side is using a different version of X.509” explained away interop problems

QC Examples: The Trivial

Software crashes when it encounters a Unicode or UTF-8 string (Netscape)

- Some other software uses Unicode for any non-ASCII characters, guaranteeing a crash
- At least one digital signature law requires the (unnecessary) use of Unicode for a mandatory certificate field
 - Standards committee must have had MS stockholders on it

Software produces negative numeric values because the implementers forgot about the sign bit (Microsoft and a few others)

- Everyone changed their code to be bug-compatible with MS

QC Examples: The Trivial (ctd)

CAs / PKI apps get subjectKeyID / authKeyID wrong (too many to list)

- CA copies subjKID into authKID field
 - Fields have a completely different structure
 - Undetected by Eudora, Mulberry, Netscape 4.x – 6.x, OpenSSL, OS X Mail, Windows
- Major CA stores binary garbage as authKID
 - No-one noticed
- European national CA encodes empty authKID

```
666   9:          SEQUENCE {
668   3:          OBJECT IDENTIFIER authKeyID
673   2:          OCTET STRING, encapsulates {
675   0:          SEQUENCE {}
      :          }
```

QC Examples: The Trivial (ctd)

- CAs create circular references for authKID / subjKID
 - AIA / altNames can also contain circular references (URLs)
 - “Processing” this extension presumably requires an infinite loop
- Not a big problem, most apps seem to ignore these values anyway (obviously)

The other CA didn't populate the [field] at all, justifying it with “Everything ignores those anyway, so it doesn't matter what you put in there”

— Peter Gutmann on ietf-pkix

QC Examples: The Serious

Known extensions marked critical are rejected; unknown extensions marked critical are accepted (Microsoft)

- Due to a reversed flag in the MS certificate handling software
- Other vendors and CAs broke their certificates in order to be bug-compatible with MS
- Later certs were broken in order to be bug-compatible with the earlier ones

Software hard-codes the certificate policy so that any policy is treated as if it was the Verisign one (Microsoft)

- Some implementations hardcode checks for obscure cert constraints
- c.f. Dhrystone detectors in compilers

QC Examples: The Scary

CA flag in certificates is ignored (Microsoft, Konqueror/
KDE, Lynx, Baltimore's S/MIME plugin, various
others)

- Anyone can act as a CA
- *You* (or Honest Al down at the diner) can issue Verisign certificates
- This was known among PKI developers for *five years* before widespread publicity forced a fix

CA certs have basicConstraints CA = false (Several large
CAs, PKIX RFC (!!))

- No-one noticed

QC Examples: The Scary (ctd)

Survey of CA certs in MSIE by Laurence Lundblade
found:

- 34 had basicConstraints present and critical
- 28 had basicConstraints present and not critical
- 40 did not have basicConstraints present
 - Some of these were X.509v1

So have CAs also issued EE certs with basicConstraints
CA = true?

- Yes
 - Consider the interaction of this with the implicit universal cross-certification model

QC Examples: The Scary (ctd)

Toxic co-dependency of broken certs and broken implementations

- Programmer has a pile of broken certs from big-name CAs/the PKIX RFC
- Ignoring basicConstraints makes them “work”
- CAs can continue issuing broken certs; implementations can continue ignoring basicConstraints

There is a fine line between tolerant and oblivious. A lot of security software which is built around highly complex concepts like PKI works mostly because it's the latter.

— Peter Gutmann on ietf-pkix

QC Examples: The Scary (ctd)

Software ignores the key usage flags and uses the first cert it finds for the purpose it needs (a who's who of PKI vendors)

- If Windows users have separate encryption and signing certs, the software will grab the first one it finds and use it for both purposes
 - This makes things less confusing for users

QC Examples: The Scary (ctd)

- CryptoAPI ignores usage constraints on keys for user convenience purposes
 - AT_KEYEXCHANGE keys (with corresponding certificates) can be used for signing and signature verification without any trouble

When I use our CSP to logon to a Windows 2000 domain, the functions SignHash AND ImportKey are both called with the AT_EXCHANGE !! Key. The certificates [...] only requires the keyusage DS bit to be true. So the public key in the certificate can only be used to verify a signature. But again: [...] the key is also used to Import a Session key. This is NOT allowed because the keyusage keyenc is not defined.

— Erik Veer on the CryptoAPI list

QC Examples: The Scary (ctd)

- Large PKI vendor ran an interop test server
 - Successfully tested against a who's who of other PKI vendors
 - After 2 years of operation, I pointed out that the certs' critical key usage didn't allow this
- European govt. organisation marked signature keys as encryption-only
 - No-one noticed
- European CA marked signature key as non-signature key
- Another CA marked their root cert as invalid for cert signing
 - Other CAs mark keys as invalid for their intended (or any) usage
- CA reversed bits in keyUsage

QC Examples: The Scary (ctd)

- The self-invalidating cert
 - Policy text: Must be used strictly as specified in keyUsage
 - Key usage: keyAgreement (for an RSA key)

What happens when you force the issue with sig-only algo?

I did interop testing with outlook, netscape mail, and outlook with entrust s/mime plugin [...] at that time I could elicit a blue screen and crypto library internal error from outlook and netscape mail respectively by giving them a DSA cert (marked with key usage of sig only). (How I came to this discovery was I tried imposing key usage restrictions and they were ignoring key usage = sign only on RSA keys, encrypting to them anyway, so I figured well let's see them try to encrypt with DSA, and lo they actually did try and boom!)

— PKI app developer

QC Examples: The Scary (ctd)

Hi. My name is Peter and I have a keyUsage problem. Initially it was just small things, a little digitalSignature after lunch, maybe a dataEncipherment after dinner and keyAgreement as a nightcap. Then I started combining digitalSignature and keyEncipherment in the same certificate. It just got worse and worse. In the end I was experimenting with mixing digitalSignature and nonRepudiation, and even freebasing keyCertSigns. One morning I woke up in bed next to a giant lizard wearing a Mozilla t-shirt, and knew I had a problem.

It's now been six weeks since my last nonRepudiation...

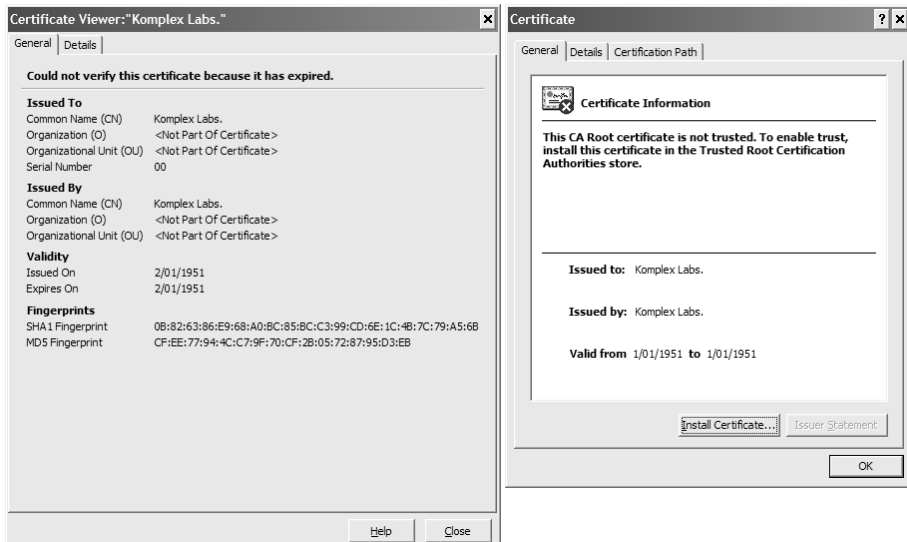
— Peter Gutmann on ietx-pkix

QC Examples: The Scary (ctd)

Obviously bogus certificates are accepted as valid (MS)

```
-----BEGIN CERTIFICATE-----
MIIQjQCCICoCAQAwDQYJKoZIhvcNAQEBBQAwGDEWMBQGA1UEAxMNS29tcGxleCEM
YWNzLjAeFw01MTAxMDEwMDAwMDBaFw01MDEyMzE1aU5NT1laMBgxFjAUBGNVBAMT
DUTvbXBsZS2XggTGficy4wggggMA0GCSqGSIb3DQEBAQUAA4IIIDQAwgggIAoIIAQCA
A+++++HELLO+THERE++++
+//And/welcome/to/the/base64/coded/x509/pem/certificate/of/+
+//KOMPLEX/MEDIA/LABS/+
+//www/dot/komplex/dot/org/+
+//created/by/Markku/Juhani/Saarinen/+
+//22/June/2000//dw3z/et/komplex/dot/org/+
+//You/are/currently/reading/the/public/RSA/modulus/+
+//of/our/root/certification/authority/certificate/+
+//Which/happens/to/be/16386/bits/long/+
+//And/fully/working/and/shit/+
+//And/totally/insecure/+
+//You/can/save/this/text/to/a/file/called/foe/dot/crt/+
+//Then/click/on/it/with/your/explorer/and/you/can/see/+
+//that/your/system/doesn't/quite/trust/the/komplex/root/+
+//CA/et/+
+//But/that's/all/right/+
+//Just/install/it/+
+//And/you're/happily/part/of/our/16386/bit/public/key/+
+//infrastructure/+
+//One/more/thing/+
+//Don't/try/read/this/with/other/PKI/or/S/MIME/software/+
```

QC Examples: The Scary (ctd)



QC Examples: The Scary (ctd)

- Validity period is actually January 1951 to December 2050
 - At one point MS software was issuing certificates in the 17th century
 - This was deliberate

the text should be changed to address the question of dates prior to 1950

— MS PKI architect on ietf-pkix

I agree with this. Every time I load one of these pre-1950 certs into the ENIAC in the basement it blows half a dozen tubes trying to handle the date, and it takes me all afternoon to replace the fried circuits. My Difference Engine handles it even more poorly, the lack of extra positions in the date field breaks the teeth of several of the gears

— Peter Gutmann, in response

QC Examples: The Scary (ctd)

- Software reports validity as 1 January 1951 to 1 January 1951, but accepts it anyway
 - It actually has a negative validity period (–1 second)
- Certificate is unsigned but cert is accepted as valid

```
30 20 30 0C 06 08 2A 86
48 86 F7 0D 02 05 05 00
04 10 A1 A1 1C 22 90 61
AF 58 8C E6 5D 40 48 BF
4D 21
```

– RSA key has exponent 1, “signing” = no-op

PGP implementations performed key validity checks after the Klima-Rosa attack

QC Examples: The Scary (ctd)

CAs issue certificates with $e = 1$

- Problem was only noticed when Mozilla was modified to detect bogus RSA keys

Both of these certs have the same problem: The RSA public key has a public exponent value that is the number 1 !! [...] I'm surprised to find certs in actual use with such a public key, especially in certs issued by well known public CAs!

— Comment on Bugzilla

- Consider the interaction of this with the universal implicit cross-certification employed in browsers
- CryptoAPI uses $e = 1$ keys as a standard (documented) technique for plaintext key export

QC Examples: The Scary (ctd)

CRL checking is broken (Microsoft)

- Hard-codes a Verisign URL for CRL checks
- Older versions of MSIE, Outlook would grope around blindly for a minute or so, then time out and continue anyway
- Some newer versions forget to perform certificate validity checks (e.g. cert expiry, CA certs) if CRL checking enabled
- If outgoing LDAP (CRL fetch) is blocked, the software hangs until it times out, then continues anyway
- Outlook 2000 doesn't check for a CRL and always reports a cert as not revoked (requires registry hack to turn on)
continues...

QC Examples: The Scary (ctd)

Today I noticed that the CRLs all have a "Next Update" date of 1/29/04, and since today is 3/26/04, I can't understand how these CRLs could still be working [...] I have been able to test that even when the "Next Update" date on CRLs has passed, IIS is still processing connection requests normally [...] Since the last post, I've been continuing to try all manner of things to try to get Windows 2000 AS to actually "care" about the validity period of the CRL, but unfortunately, have failed [...] This may be a nuance with IIS 5.0, but many applications treat no CDP in certs as an indicator that revocation does not need to be checked.

— Excerpts from a thread in MS security groups

- Outlook 2002 checks for a CRL but can't determine whether a cert is revoked or not (CRLDP-related bug)

Behaviour is representative of other PKI apps

The Lunatic Fringe

Certs from vendors like Deutsche Telekom / Telesec are so broken they would create a matter/antimatter reaction if placed in the same room as an X.509 spec

Interoperability considerations merely create uncertainty and don't serve any useful purpose. The market for digital signatures is at hand and it's possible to sell products without any interoperability

— Telesec project leader (translated)

People will buy anything as long as you tell them it's X.509

(shorter translation)

How far can you trust a PKI app?

After a decade of effort, we've almost made it to the first step beyond X.509v1 (basicConstraints)

There's not a single X.509v3 extension defined in PKIX a PKI designer can really rely on. For each and every extension somebody planning/deploying a PKI has to check each and every implementation if and how this implementation interpretes this extension. This is WEIRD!

– Michael Ströder on ietf-pkix

We're expecting banks to protect funds with this stuff?

Having worked with PKI software, I wouldn't trust it to control access to our beer fridge.

– Developer, international software company

Implementation Problem Redux

Certified for use with Windows / WHQL

- Microsoft owns the trademark
- Submit software to Microsoft, who perform extensive testing
- Passing software can use the certification mark
- Reasonable (given the size of the deployed base) interoperability among tested products
- Certified software provides certain guarantees
 - UI style
 - Install / uninstall procedures
 - Interaction with other components
 - Use of registry, correct directories, per-user data, etc etc

Implementation Problem Redux (ctd)

S/MIME

- RSADSI owns (owned) the trademark
- Simple interoperability test for signing and encryption
 - Anyone could participate, at no cost
 - Send signed + encrypted message to interop site
 - Process returned signed + encrypted message
- Passing software can use the certification mark
- Good interoperability among tested products

Implementation Problem Redux (ctd)

X.509

- No quality control
- You cannot build software so broken that it can't claim to be X.509v3

Fixing the Quality Problem

1. Create a brand (WHQL, S/MIME, ...)
2. Certify software to the brand
3. Tell users that if it has the brand, it's OK
 - (If it doesn't have the brand, it could do absolutely anything)

How not to Test

Not another industry consortium

“You've-never-heard-of-us consortium plans to have a test plan ready for X.509v7”

Not another comprehensive test suite

- Test as many obscure and rarely-used features as possible
 - Vulnerable to implementation tuning / Dhrystone detectors
- X.509 is far too complex to ever test properly
 - Follow any 2-300 message PKIX thread for examples
 - Continuous flow of new extensions and updates make all cert semantics highly mutable
 - What constitutes a pass? (nonRepudiation, anyone?)

How to Test

Just get the basics right

- Cert fetch
- Validity check
- basicConstraints / keyUsage enforcement

Simple enough that there's a single unambiguous pass / fail measure

Tests are designed to quickly catch common bugs

Lookup

App can locate the certificate it needs for an email address (S/MIME), domain name (IPsec), web server (SSL/TLS)

- Checks usability with standard Internet security protocols

App can handle multiple returned certificates

- Choose encryption cert for encryption
- Choose signing cert for signing
 - Catches lack of keyUsage enforcement

Verification

CA-issued cert contains online check URL

- CA server can be contacted at this URL

App reports valid cert † as valid

App reports invalid cert as invalid

App reports forged (manufactured) certificate as invalid

- Catches implicit universal cross-certification problems, any CA in the TCB can spoof any other CA

Verification (ctd)

App reports now-invalid cert † as invalid

- Catches the all-too-common re-read the old CRL trick
- Use blinding to detect cheating

App warns of inability to contact validation server

- Catches apps that time out and continue anyway

CA-side Cert Handling

CA cert handling

- CA cert
 - basicConstraints true
 - keyCertSign set
- EE cert
 - basicConstraints false
 - keyCertSign not set
 - digitalSignature or keyEncipherment set
 - Some CAs create lamp test keyUsage entries
 - Key is valid (e.g. no e = 1)

Catches broken CAs

Client-side Cert Handling

Client-side / application cert handling: CA certs

- basicConstraints set, keyCertSign set → accept
- basicConstraints not set or keyCertSign not set → reject
 - Catches lack of basicConstraints, keyUsage enforcement
- Rejects CA certs with invalid keys (e.g. e = 1)

Client-side / application cert handling: EE certs

- Can encrypt/decrypt with encryption cert
- Can't sign/verify with encryption-only cert
- Can sign/verify with signature cert
- Can't encrypt/decrypt with signature-only cert
 - Catches lack of basicConstraints, keyUsage enforcement
- Rejects EE certs with invalid keys (e.g. e = 1)

Challenge #6

Implementer / User Apathy (HCI)

Users find PKI incomprehensible

Why does X.509 do otherwise straightforward things in such a weird way?

[The] standards have been written by little green monsters from outer space in order to confuse normal human beings and prepare them for the big invasion
— comp.std.internat

- Someone tried to explain public-key-based authentication to aliens. Their universal translators were broken and they had to gesture a lot
- They were created by the e-commerce division of the Ministry of Silly Walks

Consequences of lack of user understanding

PKI-enabling an app is just a side-job for developers

- Motivation: The boss said do it
I don't need to pay Verisign a million bucks a year for keys that expire and expire. I just need to turn off the friggen [browser warning] messages.
 - Mark Bondurant, alt.computer.security
- Get it out of the way as quickly as possible
 - CA generates key and mails it out
 - Private key is shared across as much of the org. as possible
 - “Revocation check” repeatedly re-checks against the same old CRL stored on disk
- Meets all PKI checkbox requirements without having to go to the effort of getting it right

Default-to-Secure Design

Make the right way the only way to do it

- PnP PKI makes it very hard to not do local key generation, distinct signature and encryption keys, minimised TCB (trusted CA certs), keys generated in hardware, etc etc
- Realtime validity check makes it very hard to just go through the motions of performing the check

KISS

Simple design discourages homebrew (= insecure) mechanisms

```
cryptCreateSession session
cryptSetAttribute session, _
    CRYPT_SESSINFO_SERVER_NAME, "[Autodetect]"
cryptSetAttribute session, _
    CRYPT_SESSINFO_USERNAME, userName
cryptSetAttribute session, _
    CRYPT_SESSINFO_PASSWORD, password
cryptSetAttribute session, _
    CRYPT_SESSINFO_PRIVKEYSET, keyset
cryptSetAttribute session, _
    CRYPT_SESSINFO_ACTIVE, true
```

This is the entire PnP PKI (Challenge #2) interface

KISS (ctd)

Other operations are similarly idiot-proof

```
crypt.CreateSession( session );
status = crypt.CryptCheckCert( certificate,
    session );
```

This is the complete real-time validity checking (Challenge #3) interface

Conclusion

Certificate lookup

- Simple HTTP interface uses the web as a Public File

Enrolment

- PnP PKI eliminates enrolment pain, makes it easy to do the right thing

Certificate validity check

- Real-time online check matches requirements for online banking, etc

Quality control

- Core functionality checked through simple, unambiguous tests

Postscript: Implementation Availability

Available as the cryptlib security toolkit,

<http://www.cs.auckland.ac.nz/~pgut001/cryptlib/>

Implementation and usage details

- ANSI C, runs on anything: BeOS, DOS, eCOS, μ TRON, Mac OS X, MVS, QNX Neutrino, RTEMS, Stratus OS, Tandem NSK, Unix (any variant), Win16, Win32, WinCE/PocketPC, VxWorks, VM/CMS, no OS (runs on the bare metal)
 - Minimum RAM requirement: ~128K (may run in 64K)
 - Please contact the author if using one of the more obscure embedded/RTOS systems with special considerations
- Open-source implementation, dual-licence
 - GPL or standard commercial license, your choice

Experiences of establishing trust in a distributed system operated by mutually distrusting parties

Scott Crawford, Enterprise Management Associates, Boulder, CO, USA

Email: scrawford@enterprisemanagement.com

David Chadwick, University of Salford, Salford, England, M5 4WT

Email: d.w.chadwick@salford.ac.uk

Introduction

The organization that is the subject of this case study is engaged in the worldwide monitoring of environmental information. This information provides evidence about the production of contaminants in one country that can be harmful to its neighbors. The project, which started in early 1999, was to develop an IT system that could authenticate data collected from widely distributed sources, in a manner that could be trusted by the participating countries, even though those countries might not trust each other.

Because of the potential political implications of this monitoring activity and the data collected, the subject organization represents the interests of the governments of the participating countries. Therefore, conclusions drawn from the collected data must be based on a reasonable degree of trust in the integrity and authenticity of the data and the data collection, archiving and distribution systems involved.

Because the interests of multiple sovereign and independent nations are involved, none of the participating nations is willing to subordinate its national interest to the subject organization by allowing the organization to speak on its behalf regarding the veracity of the data or any evidentiary conclusions that may be drawn or implied. For example, while a chemical or nuclear accident such as the 1984 Bhopal, India or 1986 Chernobyl disasters could conceivably produce contaminants indicative of hostile military activity, to draw such a conclusion from the data collected from a similar accident would be in error, but not outside the realms of possibility for one nation seeking to thwart the national interest of another nation in which such an accident had taken place. Therefore, one of the principal goals of the data authentication system was to assure that the trust placed in it – and, by extension, in the data itself – be a matter shared amongst, if not all, at least a significant enough number and distribution of participant nations to give a reasonable assurance to the organization as a whole that the integrity and veracity of the data is trustworthy. Each participant nation or any other observer would then be free to draw their own conclusions as they see fit.

To support this goal, the monitoring regime involved in collecting the data was developed along the following lines:

- Its human structure paralleled that of the organization itself. Its policy-making bodies were designed to be democratic and deliberative, and its operational staff developed along lines of proportional representation of participating nations, with oversight by the representative policy-making organs.

- Technical systems were developed to reflect the collective and representative nature of the organization. The data collection system was designed around the placement and distribution of monitoring sites worldwide, with locations distributed among as many participating countries as possible to monitor the global environment as a whole. A number of scientific disciplines were involved, for purposes of confirmation and cross-referencing of data indications.

Monitoring is continuous wherever possible. Monitoring sites have been networked with data management centers to enable timely collection and analysis. The data itself, as well as the data collection, archiving and distribution systems, are open to the scrutiny of all participating nations. The influence of any one or minority of participants on the data – particularly nations hosting monitoring and networking sites – should be minimized as far as possible. This was to be achieved by the participation of representatives of the subject organization in the construction, operation and maintenance of the data collection sites.

System Trust Requirements

These principal considerations influenced the nature of the IT system that was developed to authenticate the veracity and integrity of the collected data. Organizational policy-makers required the authentication system to implement an architecture that distributed the trust among the participants. Policy-makers further required the authentication architecture to parallel the construction of the data collection system and to be open to the highest possible scrutiny and periodic evaluation by representative groups. This requirement, however, had to be balanced against the need to protect the system and its individual sites and components from exploitation. For example, a malicious party seeking to blind the organization to polluting or contaminating activity in a specific location might seek to interfere with the monitoring ability of a site through interfering with its network connectivity or system operation. It also had to be balanced against the risks posed by a pragmatic need to delegate contractual, implementation and operational responsibilities to those having the necessary expertise. Such delegation was, however, subjected wherever possible to oversight by representative groups reflective of the collective nature of the subject organization as a whole. An overriding principle was that no part of the system installation or operation that formed part of the trust infrastructure, should be entrusted to a single individual.

Proposed Solution

Because the monitoring data could be represented as either a networked bitstream or a discrete message, it was determined that digital signatures could be applied as a means of assuring data authenticity. Pioneered by [DH 76] and elaborated in [RSA 78] and subsequent innovations and standardizations (e.g. the PKCS#1 standard for the RSA algorithm [PKCS 1]), public key cryptography (PKC) implements digital signatures through the combination of public/private keypairs and hash algorithms. Encryption of the data with a private key and successful decryption with the corresponding public key assures that only a specific private key could have performed the encryption. If the encrypted data is a “one-way” hash of the actual subject data, such as provided by the Secure Hash Algorithm (SHA-1, [FIPS 180-1]), it provides a tamper-evident assurance that the data has not been altered since encryption, when the decrypted hash matches one generated over the received data. The authentication system in this application was therefore centered on digital signing of the data at the monitoring site at the time of observation, and as close to the data source as possible so as to limit opportunities for data alteration. Streaming networked data could be divided into discrete transmission frames to which individual digital signatures could be applied.

But while digital signatures may provide a mechanism for authenticating data, of themselves, they do not address the issue of distributing trust amongst the participating nations. For example, a recipient needs to know which private keys have been installed at which data monitoring sites, and that it has the correct corresponding public keys in its possession.

Review of Previous Work

Prior to implementation, other relevant implementations of PKI technology were studied so as to provide insight into how trust can be distributed amongst competing, and possibly mutually distrusting, member organizations. Candidate organizations were multinational organizations involved in international finance, banking and exchange systems. The stakes of individual participants in these multinational organizations, concerning the authenticity of information and data exchange, which represents large sums of money, are at least as significant as the risks borne by the participants in the subject organization.

One of the most significant parallels was found in the establishment of Identrus LLC, which took place at approximately the same time as the early stages of the subject implementation. Founded in April 1999 by eight leading US and European banking and financial institutions, Identrus was created for the purpose of establishing an architecture of trust in electronic transactions between participating banks and institutions, and between their customer businesses as well ([Identrus 98], [Identrus 02a]). One of the original participants in Identrus was the US company CertCo ([Identrus 98]). CertCo was differentiated from its competitors at the time by its implementation (with IBM cryptographers) of threshold public key cryptography ([Ankney 00]). [Desmedt 92] describes the goal of threshold PKC as a scheme “in which the power to perform a certain operation is shared.” In a threshold cryptosystem, the factors of a key are *distributed* among a group such that, when the group members contribute their factor components for combination enabling an encryption operation, they do so without divulging their individual components to each other. More to the point, a threshold cryptosystem requires a minimum *threshold* number m out of the total number n (described as “ t -out-of- l ” in [Desmedt 92]) of all possible participants to contribute their components in order to enable the encryption operation.

Threshold cryptography was therefore studied as a possible enabling technology for the distribution of trust between the cooperative yet mutually-distrusting participants in the subject organization. However, a threshold technique posed significant operational complications when considered for application of digital signatures at the data source of an environmental monitoring station. Instead, attention turned to a threshold implementation in the management of the data-signing keys. Because a system of digital signatures relies on the integrity of the private keys used to generate the signatures, a system of management of the corresponding public keys predicated on the then-current X.509v3 standard of digital certificates [X.509] was decided upon. The use of threshold cryptography in generating the digital signatures on the certificates of the issuing certificate authority (CA) was considered.

Threshold cryptography was not, however, a panacea without its own flaws. [Langford 96] illustrated certain vulnerabilities in systems then current: A colluding subgroup of the minimum required number of participants was able to manipulate a forgery of a threshold signature without the knowledge of the other participants (effectively reducing m to 2-out-of- n , regardless of the intended size of m). A malicious participant was able to influence public key generation such that they were enabled to discover the complete private key which is supposed to be unable to be discovered by any participant or used without the threshold number of participants. The conclusion drawn by [Langford 96] was that systems “without a trusted key generation center...are more complicated than those that do allow a single trusted center and are therefore more vulnerable to manipulation.” [Desmedt 97] pointed out that not all threshold algorithms had progressed to an equal state of security in their development. In particular, [Desmedt 97] noted that, at the time, “no practical threshold [implementation of] DSS [the

Digital Signature Standard implementation of the Digital Signature Algorithm (DSA), now [FIPS 186-2]]...has been presented so far.”

The lack of threshold DSS conflicted directly with the preference of a number of the participant nations for the use of DSS in data signatures. At the time, many national governments had concerns regarding the export of encryption technology, and did not want an organization representing their national interests to be accountable for potentially enabling undesired access, potentially worldwide, to an encryption technology such as RSA, in which either public or private keys could be used to encrypt digital information. DSS was therefore preferred, as DSS private keys could be used (in principle) *only* for signature generation, and the corresponding public keys *only* for signature verification. Thus, non-DSA-based algorithms – including threshold cryptosystems then available – were ruled out.

The example of distributed trust as manifested in *m*-out-of-*n* threshold key management and certificate authority implementations was, however, retained in a requirement to implement a distributed key management system. A “mixed” system of threshold-based certificates of DSS signing keys was briefly considered, but abandoned due to the above-related issues with threshold cryptography and algorithm preference, as well as the problems foreseen for a system of mixed algorithms. Instead, an administrative, rather than technical, implementation of *m*-out-of-*n* signature generation in the issuance of DSA-signed certificates was undertaken. After an evaluation of solution providers worldwide, the UK company Baltimore Technologies was selected to provide tools and systems for implementing an *m*-out-of-*n* key management system predicated on DSA. A number of other vendors from several different nations participated in the implementation of DSA signature generation software at the data sources.

Initial implementation

The Baltimore Technologies implementation was selected, in part, because of its flexibility in “customizing” a CA architecture to the needs of an organization, including its ability to use multiple Registration Authorities (RAs) and Registration Authority Operators (RAOs) to meet the administrative *m*-out-of-*n* requirement in the issuance of digital certificates. In Baltimore’s PKI, RAs are client systems that submit requests for an X.509v3 digital certificate to an issuing CA server. RAOs are parties (usually humans) that interact with the RA to enable the approval of a certificate request for forwarding by the RA client to the CA. The certificate request comprises the public key to be certified and other relevant information about the key and its holder, formatted according to the PKCS#10 standard [PKCS 10]. Baltimore’s PKI supports both single and multiple RAs interacting with a CA to request a certificate, as well as multiple RAOs interacting with an RA before a request can be sent to the CA. By mandating that multiple RAOs must request the same certificate to be issued for a data monitoring station, effectively distributes the trust placed in the operation of the CA to the number of RAOs that are involved in issuing the certificate requests.

The implementation was staged over periods of preliminary design, pilot testing, final design prior to initial implementation, and the initial implementation itself. Laboratory implementations of the data signing architecture were developed to test the processes of: keypair generation, certificate request and issuance involving *m*-out-of-*n* RAOs, signature of actual data, transmission of data and signatures via networks, management and retrieval of digital certificates, and the use of certificates in signature verification of data. Parameters and issues of general system operation and maintenance were also evaluated.

As preliminary system design took shape, distribution of trust in the system became manifested in a variety of ways beyond key management *per se*. As noted earlier, a number of scientific disciplines were involved in the monitoring regime, to give corroboration and cross-referencing of data supporting indications of specific contaminants and contaminating actions. Thus trust was distributed across a number of monitoring techniques, from measurement of atmospheric compounds to highly sensitive detection of vibrational information transmitted through the earth's oceans and the earth itself. Such multiplicity of data sources and types contribute to the weight of evidence in any given case, even in cases where signature-based authentication at any one monitoring site or minority of sites might be compromised.

Multiple parties were also involved in the construction and deployment of specific monitoring sites as well as the central data collection and management points supporting the system as a whole. In each case, representatives of the entire range of participating nations were involved, reducing the possibility of subversion of critical system components at virtually every key point.

Distributing Responsibility

Such a distribution of responsibility was not, however, without its cost. In the development of the authentication system, at least six, and sometimes more, different contractors spread throughout the world were involved in the detailed technical specification of the various components of authentication. In some cases, different contractors were delegated responsibility for the elaboration of the signature-implementation systems for different monitoring disciplines. Differences in standards were also required for different data transmission techniques (i.e. networked bitstreams versus discrete or "segmented" messages). The organization defined its own standard technique for signing streamed data by allowing a 40-byte space for a DSS signature in each transmission frame. Segmented message-format data had to be signed according to a standard that could be interpreted by both the implementing contractors and the subject organization. The standard chosen was that in most common use at the time, S/MIMEv2 (Secure Multipurpose Internet Mail Extensions) [RFC 2311]. S/MIME defines how to create a MIME body part that has been cryptographically protected according to [PKCS 7]. However, neither S/MIME nor PKCS#7 define the object identifier to be used with the DSA/DSS signing algorithm (they only specify ones for use with RSA). Therefore, accommodation was required among the contractors to enable the DSS signing algorithm. Laboratory implementations were ultimately successful using a variety of tools, including adaptation of open source reference implementations such as OpenSSL (then at version 0.9.5).

One of the implications of the unique nature of the monitoring regime was the necessity for custom developments in certain monitoring installations. For example, certain subterranean monitoring installations at deep levels below the earth's surface posed special problems for system endurance and form factor, as did underwater detectors placed beneath the ocean's surface. In certain cases, placing signature-generation devices at the *exact* point of data collection were impractical. In many cases significant barriers and challenges had to be overcome. For example, in some installations, the technologies necessary to compose standard certificate requests strictly formatted to PKCS#10 were beyond the physical and technological constraints of the systems at their then-current state of development. The solution to this involved on-site personnel obtaining "raw" public key information from such devices. The absence of a formal PKCS#10 request and an associated signature generated by the corresponding private key (which, when verified by the public key contained within the

PKCS#10 request, is a crucial step in demonstrating certificate request integrity and private key ownership) would be compensated for by the presence of on-site observers who received and verified the integrity of the generated public key from the remote constrained detector. The public key material thus obtained would then be sent in one or more PKCS#7-compliant messages to a key management center by the on-site observers. At the key management center, an adaptation of certificate issuance systems was developed to permit direct submission of such public key material to the CA when verified by the RAOs. Laboratory tests of this combination of techniques were successful in obtaining a certificate for a keypair generated in this manner. Demonstration of the integrity of the process was verified by the auditable recording of participant actions in order to preserve the “chain of trust”.

This technique illustrates one example of how the presence and participation of multiple persons at virtually every crucial step of the authentication system became essential to establishing and maintaining the concept of distribution of trust, necessary for the system as a whole. Implicit in such a system, however, is the necessity of informed human participation; but this, after all, is to be expected in a system predicated on trust, which is essentially a human phenomenon. A certificate-authority-based key management architecture is, by definition, based on an assumption of trust in the authority itself. Trust, however, may be interpreted and manifested in a multiplicity of ways ([Mayer et al 95], [AJ 98], [Kramer 99]). Multiple parties may not – perhaps *will* not – all agree on their individual perceptions of what is trustworthy and what is not. However, the assumptions made in the design of this system considered that when a significant number of participants were agreed that they could place their trust in a system consisting of a number of verifiable measures and components, the requirement for trust distribution would be satisfied.

This also, however, implied that a certain number of duplications in implementation would be necessary to assure the necessary participation of multiple parties at significant points in the architecture. No one person could be allowed to operate alone in the presence of crucial system components, when those components might be susceptible to exploitation by an individual. The system would have to enforce multiple authorizations for access, manipulation and control beyond the requirement of *m-out-of-n* necessary for certificate issuance. The possibility existed that system operators might be required to be responsible for several components such as smartcards and other tokens necessary to enable operation of certain system elements. Backups of key material would have to be distributed among a number of points, all in an auditable fashion.

To meet these exigencies, a minimum set of qualities were sought as design goals. The threshold number of persons or components among which crucial elements of the system would be distributed would be kept to as practical a minimum as possible without subjecting the system to the susceptibility of individual operators. This did not rule out the actions of a malicious minority in all cases, but the sheer preponderance of numbers of persons and steps toward authentication involved throughout the architecture mitigated the possibility of such isolated actions subverting the system as a whole. Standards of procedure and operation would also be developed, with the intent that persons interacting with the system would be informed and knowledgeable. Operators would be instructed regarding what they would be doing and the reasons why trust in the system would be enabled by their actions, while those depending on the system for trusted demonstrations of authenticity would be aware of how and why the system should be trusted. System operations as well as the signed data itself would be auditable and open to scrutiny by appropriate parties, thus fostering the openness necessary to the development

of trust described in [Mayer et al 95], [AJ 98], [Kramer 99], and others. The use of OpenSSL in bespoke development of authentication components, for example, enabled clear examination of source code used in implementing authentication. Wherever possible, similar cooperation was obtained from contractors, sometimes in the form of “source code escrow,” preserving the contractor’s proprietary rights in maintaining source code confidentiality while enabling the organization to have the option of source code review should it be desired.

It would be inevitable that, beyond outright exploit, human as well as technical errors would eventually begin to be manifested in such a system, perhaps posing a more significant threat than malicious exploit. Again, however, the preponderance of the number of monitoring sites, the numbers of points throughout the architecture in which multiple parties would be involved, and the numbers of persons involved in critical operations, mitigated the potential consequences of any one error or a small number of errors. Added to these factors are data management systems at the data centers receiving the signed data that are able to alert operators when signature verification failures occur. The data centers hosted by individual participating nations help to verify the validity of such incidents and may themselves track such occurrences independently, thus helping to keep them from being hidden in a possible exploit scenario. Thus, a general development of “trust by consensus” in which the number of individual actions and steps in data authentication accumulate towards a body of data supporting trust, began to emerge as the system design progressed.

In summary, a description of the initial implementation proposed for the distributed management of trust is as follows. At the time a monitoring site is to be enabled with a digital signature capability, an on-site team of operators generates the keypair. If satisfied with the key generation process and the integrity of the resulting keypair, the on-site team then forwards the resulting public key to the key management center in one or more PKCS#7-compliant messages bearing the digital signatures of the on-site observers. At the key management center, the signatures of the received messages are verified against the signer’s certificate(s) by a group of authorized RAOs. If a minimum m out of a total number of n RAOs agree that the signed message(s) containing the submitted public key are trustworthy based on signature verification and other verifications of the on-site observers’ presence at the site, the RAOs approve the certificate request, and the certificate is duly issued by the CA. Signed data thereafter received from the site is verified on receipt at the data management center by signature verification using the issued certificate accessed from a local directory of certificates and certificate revocation lists (CRLs are as defined in edition 3 of X.509 [X.509]). This directory is accessed according to the Lightweight Directory Access Protocol (LDAP) described in [RFC 1777] and [RFC 2251], with an organizational namespace rooted on the name of the organization itself (being, as it is, an international entity).

System Maintenance

Yet to be fully elaborated are issues of key rollover and replacement of valid keys. The assumption to date has been that as the current key lifetimes reach their pre-determined limit (set to a minimum of 5 years) PKC technology will have matured to the point where a more evolved implementation may be indicated. Regardless, a preliminary protocol has been worked out, in which a currently trusted signing key is used to “countersign” the certificate request generated for a new keypair. Thus, a data generating system needs to be able to “cache” a currently valid keypair while awaiting issuance of the replacement keypair’s certificate and authorization to use the new signing key. In cases where such caching is not possible, data would need to be signed

immediately with the new signing key. Authentication by the recipient is then contingent upon the issuance of the replacement certificate for the new keypair, during which time the validity of the site's data would be in a state of suspension. In any event, under such a scheme a new keypair would not be generated except on the site's receipt of an authenticated (digitally-signed) command issued by a minimum number of authorized parties, identified by certificates available to the site itself. Unauthorized keypair generation messages would be detected when data signed by an unauthorized key is received at the data management center. No authorized certificate would be available for data verification, and authentication would fail. In this case data from such a site would be "suspended" from authentication until on-site remediation was undertaken to restore the site to a trusted state.

In such a scenario, monitoring sites would need to have certificates of authorized command-issuers available to them, in order to enable command authentication. For any site installation, a certain number of individuals and groups must be delegated the responsibility for trusted operations on the site. Again, the limits of trust related to the number of individuals authorized to operate at a site is mitigated by the numbers of sites, the numbers of persons involved, and the oversight of such operations made possible by open scrutiny and the auditability of actions. In its role as the facilitating entity for the regime as a whole, the subject organization is able to call on on-site representatives and other means to monitor and corroborate site changes. It is also able to track any site changes that are authorized, thereby further mitigating the risks arising from trust placed in the site.

At the site, authorized signed commands are distinguished by the positive identification of the command-issuer through verification of the issuer's digital certificate. Without such verification, commands are ignored. To maintain the availability of the most up-to-date certificates, as well as information regarding suspended or revoked certificates, two methodologies were proposed. One was network-based access to directories containing certificates and CRLs; the other was on-site storage and maintenance of the necessary certificates and CRLs. The ultimate goal in the future is network enablement of certificate status checking for the most timely validation by the monitoring sites, using a protocol such as OCSP ([RFC 2560]). However, not all sites are currently capable of supporting such technological demands. On-site storage of current necessary certificates and CRLs will supplement such sites. In such cases, it is possible that, for example, an authorized individual whose authorization has been revoked may command a site as yet unaware of the revocation. Such cases are mitigated by limitations on access to the command-and-control functionality itself, and by requiring more than one individual to issue the same command (but this latter functionality has not been implemented yet). Also, the number of sites and persons involved spreads the individual risks. In no case would any one individual be authorized to command more than a significant minority of sites.

Local vulnerabilities in the monitoring sites are mitigated through a number of measures intended to develop tamper-evident installations that generate alerts whenever a site exploit is attempted. This is enabled through the triggering of functionality that includes site ill-state-of-health information in the data flow indicating that the site has been accessed. Cutting off the site's network connectivity in order to "blind" either the subject organization to an exploit or the site to the existence of, e.g., revoked command-issuer certificates, is detected by the absence of expected data from the site. This data cannot be mimicked to obscure the exploit without the attacker having access to the signing keys. Replay of valid data is not an option either, since the data contains replay detection information. Even scheduled maintenance may produce alerts of site intrusion, but such alerts are verified as scheduled maintenance from published operational

plans. There are, of course, limitations to the amount of trust that can be placed in such measures, but the “trust horizon” relative to the number of persons involved and their motivation to exploit is limited to the extent practical to the maintenance and purpose of the organization itself.

Initial results of implementation

At the time of writing, approximately one-fourth to one-third of monitoring sites have been equipped with the initial implementation of digital signatures and are sending authenticated data to the data management centers. While authentication of data in these cases has been successful, some of the most significant results are as follows.

The human interaction necessary to enable the system operation described above has been considerable. In particular, one of the author’s personal experience in orienting operational staff and users to the system indicates that the learning curve alone is significant. Simply orienting users to the nature and operation of public key cryptography has been an abstraction difficult to communicate in many cases. Compensating for such challenges has been the high motivation and dedication of the subject organization’s participants to fully understand the system. Thus, it is our subjective opinion that motivation is a significant factor in the success of such a trust-based system. In addition the aptitude of users to understand the operation of public key cryptography, as well as the ability of system developers to communicate the information essential to understanding, are essential to success.

The burden authentication technology places on data management systems themselves should not be overlooked. Performance measures of the time and resources necessary to authenticate a large number of DSS signatures received from stations continuously transmitting proprietary data protocol frames is not insignificant. System capacity planning and development is still taking shape to accommodate such demands. Computational and network resources are not the only demands placed on an organization seeking to implement a system such as this. Measures necessary to assure the security of all aspects of the implementation also take a toll on both human and material resources. A higher degree of vigilance and standardization of operational policies and procedures is necessary to assure the integrity of the system itself.

Nevertheless, the general consensus among users at this point appears to be divided between those who feel they understand the authentication system and those who do not. Surveys of users are currently underway to establish quantitative measures of success or failure of the implementation. Until they are received and analyzed, interviews with current system participants indicate that those who manifest an understanding of the system are satisfied that the system is functioning successfully. Nevertheless, they are not happy with the burden imposed by both the procedural and computational requirements of this distributed-trust implementation of digital-signature-based authentication. They are also less than satisfied with the “usability” of the human-interactive components of the system, which can be cumbersome and require the necessary understandings which can be challenging to communicate effectively to the involved personnel, as described above. Among those who do not express a high understanding of the system, the above-described burdens appear to be regarded as excessive relative to the benefits that are derived. It is not yet clear whether further education and dissemination of information regarding the system and its necessity to the requirements of the organization would help alleviate such concerns. However, future developments in the system will almost certainly take such measures into account.

Summary and conclusions

The most obvious conclusion drawn from this experience is that “distributed trust” means, first and foremost, distribution among people. While that statement may seem so apparent as not to even require being made, consider that first- and second-generation public key infrastructures (PKIs) such as this almost universally began as technological exercises. Focus on the algorithms and the technology of encryption and digital signature led to a number of implementations which did not sufficiently consider the human factors of trust, as well as other human factors such as how people perceive technology, how such perceptions affect their use of technology, and the relationship of such perceptions to assumptions – correct or not – made by system planners and developers, particularly as affects the success or failure of security technologies. This, in turn, led to the development of a framework for certification practice and certificate policies such as that described in [RFC 2527], but policies alone are not enough to compensate for the involvement of human beings in the technology of trust. Studies of human factors in security implementations such as [WT 98] and [WT 99] demonstrate that what often begins as a technology exercise often ends, successfully or not, as an exercise in implementing an appropriate understanding of the human factors involved. It is our belief that the technological limits of security implementation are often subject to the fact that both security and trust are fundamentally human concepts.

In the case of this implementation, the key goal was to implement a system in which all participants could trust the distributed data as far as it was possible and practical, notwithstanding the political nature of the organization and the lack of trust between the participants. It was essential to prevent a minority of persons with malicious intent from being able to subvert or exploit the data authentication system. Initially, this effort focused on the technology of distributing trust as manifested in the threshold cryptography system of digital signatures. Ultimately, the system has become one where trust is dependent upon the sheer numbers of people and systems that are involved, viz.: the number of points of data collection and their worldwide distribution necessary to obtain a reasonable number of overlapping systems and techniques of measurement; the numbers of participating nations and their representatives; the numbers of individuals involved in critical steps in the authentication process; and the volume of data itself. It must be noted that each of these factors was in existence in this organization before the authentication system itself was undertaken. Therefore, it would seem that authentication is dependent on the existence of the underlying factors that enable the necessary distribution of trust; the authentication system itself does not enable or distribute trust independent of these pre-existing factors.

Nevertheless, the system may at this point be judged a qualified success, in so far as it has succeeded in enabling a tangible measure of trust in the authenticity of the signed data, through the participation of a number of capable systems and motivated, knowledgeable individuals. However, the organization is distinguished as one which attracts individuals from throughout the world who are motivated to see it succeed. While this may in many respects be true of most professional organizations, the subject organization is able to call on the resources of national governments owing to the political nature of its existence. While environmental monitoring may not be a high priority with many participating governments, it nevertheless distinguishes the organization from, for example, those in the private sector with more limited resources. Only those organizations not just capable of, but also having a mandate for, fielding the necessary resources in terms of motivated, knowledgeable staff and technological capacity and development would likely be interested in such an undertaking. It is therefore not surprising

that architectures for building trust into distributed systems run by mutually distrusting or wary parties have to date only been undertaken primarily by banks, international financial institutions and other entities operating in the arena of marketing top-level trust assurance. More recently, international military coalitions [FRD 02] have also been shown to have the resources and mission to do so. This may be at least partly attributable to the potentially cumbersome nature of X.509-based hierarchical PKIs and their related technologies. Developments such as SPKI [IETF 01], authorization-based certificates and “federated” trust architectures such as the Internet2 Shibboleth project [I2 03] may succeed in helping to shape trust technology more closely to the realities of human use and interaction, particularly in more common, less well-endowed environments, but as yet it is too early to say so conclusively.

As a final note, the authors wish to state that the content of this paper represents the authors’ own views. The authors do not represent or speak for or on behalf of the subject organization, nor should any statement in this paper be so construed.

References

- [AJ 98]: P. J. Ambrose, G. J. Johnson. A Trust Based Model of Buying Behavior in Electronic Retailing. Association for Information Systems 1998 Americas Conference Proceedings, pp. 263-265. In <http://www.isworld.org/ais.ac.98/proceedings/track06/ambrose.pdf>.
- [Ankney 00]: R. C. Ankney. Certificate management standards. CertCo, Inc., 2000, in <http://www.certco.com/pdf/cms.pdf>.
- [Desmedt 92]: Y. Desmedt. Threshold cryptosystems. In J. Seberry, Y. Zheng, eds. AUSCRYPT '92. Lecture Notes in Computer Science 718. Springer-Verlag, Berlin/Heidelberg/New York, 1993, pp. 3-14.
- [DH 76]: W. Diffie, M. E. Hellman. New directions in cryptography. IEEE Transactions on Information Theory, IT-22:644–654, 1976.
- [FIPS 180-1]: National Institute of Standards and Technology, US Department of Commerce. Secure Hash Standard. 17 April 1995, in <http://csrc.nist.gov/publications/fips/fips180-1/fips180-1.pdf>.
- [FIPS 186-2]: National Institute of Standards and Technology, US Department of Commerce. Digital Signature Standard. 27 January 2000, in <http://csrc.nist.gov/publications/fips/fips186-2/fips186-2.pdf>.
- [FRD 02]: G. Fink, S. Raiszadeh, T. Dean. Experiences establishing an international coalition public key infrastructure. 1st International PKI Research Workshop – Proceedings. April 2002, pp. 193-206, in <http://www.cs.dartmouth.edu/~pki02/Fink/paper.pdf>.
- [I2 03]: Internet2 Consortium. Shibboleth project. July 2003, in <http://shibboleth.internet2.edu>.
- [Identrus 98]: Identrus LLC press release. Major financial institutions announce new company to provide businesses globally with a single electronic identity. 21 October 1998, in http://204.141.53.14/knowledge/releases/us/release_102198.xml.
- [Identrus 02a]: Identrus LLC. About Identrus. 2002, in <http://www.identrus.com/community/index.xml>.
- [IETF 01]: Internet Engineering Task Force. Simple Public Key Infrastructure (spki). IETF Charter, 16 January 2001, in <http://www.ietf.org/html.charters/spki-charter.html>.
- [Kramer 99]: R. M. Kramer. Trust and distrust in organizations: Emerging perspectives, enduring questions. Annual Review of Psychology, vol. 50, pp. 569-598. Annual Reviews, Palo Alto, CA, 1999.
- [Langford 96]: S. K. Langford. Weaknesses in some threshold cryptosystems. In N. Koblitz, ed. CRYPTO '96. Lecture Notes in Computer Science 1109. Springer-Verlag, Berlin/Heidelberg/New York, 1996, pp. 74-82.
- [Mayer et al 95]: R. C. Mayer, J. H. Davis, F. D. Schoorman. An integrative model of organizational trust. Academy of Management Review. Vol. 20 no. 3, 1995, pp. 709-734.
- [PKCS 1]: RSA Laboratories, Inc. PKCS #1 v2.1: RSA cryptography standard. 14 June 2002, in <ftp://ftp.rsasecurity.com/pub/pkcs/pkcs-1/pkcs-1v2-1.doc>.
- [PKCS 7]: B. Kaliski. PKCS #7: Cryptographic message syntax version 1.5. Request for Comments 2315. IETF Network Working Group, March 1998, in <http://www.ietf.org/rfc/rfc2315.txt>.

[PKCS 10]: M. Nystrom, B. Kaliski. PKCS #10: Certification request syntax specification version 1.7. Request for Comments 2986. IETF Network Working Group, November 2000, in <http://www.ietf.org/rfc/rfc2986.txt>.

[RFC 2311]: S. Dusse et al. S/MIME Version 2 Message Specification. Request for Comments 2311. IETF Network Working Group, March 1998, in <http://www.ietf.org/rfc/rfc2311.txt>.

[RFC 2527]: S. Chokhani, W. Ford. Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework. Request for Comments 2527. IETF Network Working Group, March 1999, in <http://www.ietf.org/rfc/rfc2527.txt>.

[RFC 2560] M. Myers, R. Ankney, A. Malpani, S. Galperin, C. Adams. X.509 Internet Public Key Infrastructure: Online Certificate Status Protocol – OCSP. Request for Comments 2560. IETF Network Working Group, June 1999, in <http://www.ietf.org/rfc/rfc2527.txt>.


[RSA 78]: R. Rivest, A. Shamir, L. Adelman. A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. Communications of the ACM, 21 (2), pp. 120-126, February 1978.

[WT 98]: A. Whitten, J. D. Tygar. Usability of Security: A Case Study. Carnegie Mellon University School of Computer Science Technical Report CMU-CS-98-155, December 1998, in <http://reports-archive.adm.cs.cmu.edu/anon/1998/CMU-CS-98-155.pdf>.

[WT 99]: A. Whitten, J. D. Tygar. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. Proceedings of the 8th USENIX Security Symposium, August 1999, in <http://www.usenix.org/publications/library/proceedings/sec99/whitten.html>.

[X.500]: ITU-T (formerly CCITT) Recommendation X.500 / ISO/IEC/ITU Standard 9594 (series). Information Technology – Open Systems Interconnection – The Directory. 1988.

[X.509]: ITU-T (formerly CCITT) Recommendation X.509 (1997 E) / ISO/IEC/ITU Standard 9594-8. Information Technology – Open Systems Interconnection – The Directory, Part 8: Authentication Framework. June 1997.



Experiences of establishing trust in a distributed system operated by mutually distrusting parties

Scott Crawford

Enterprise Management
Associates
Boulder, CO

David Chadwick

University of Salford
Salford, UK



Introduction

- ⌘ Subject organization: Engaged in worldwide environmental monitoring
- ⌘ Observed contaminants/events could result from accidents as well as events of international political significance
 - ⌘ Examples: Bhopal 1984, Chernobyl 1986
- ⌘ Detection could have an impact on international relations



Introduction

- ⌘ Participants therefore represent interests of participating nations who don't always fully trust each other
- ⌘ Monitoring data must therefore **assure integrity and trust**
- ⌘ But, it must be trustworthy to organization as a whole, not too greatly influenced by any one party or a minority



Organizational structure

- ⌘ Representative policy-making bodies, with policy-making working groups focusing on specific aspects
- ⌘ Operational organization and staff representative of participants
- ⌘ Distribute responsibility for implementation among representative groups; limit influence of individual participants or hosting nations where possible
- ⌘ Thus, distribution of trust in implementation on many levels



Data collection and monitoring regime

- ⌘ Data collected from over 300 sites worldwide in multiple scientific disciplines
- ⌘ "Continuous" (waveform) and "segmented" data (discrete messages)
- ⌘ Sites networked to central international data collection facility
 - ⌘ Raw data
 - ⌘ Data analysis services also provided at central site: supplementary information, not conclusion-drawing
- ⌘ Participants' national data repositories also collecting some or all data or data subset(s)



Principles of distributed trust in data collection

- ⌘ Organization's operations division collects and provides data impartially, allowing each participant to draw their own conclusions
- ⌘ Principles:
 - ⌘ Data, operations open to scrutiny
 - ⌘ Distribute trusted responsibilities among a group
 - ⌘ No individual should be a single point of trust
 - ⌘ "Preponderance of data" from multiple collection points, disciplines



Data authentication

- ⌘ Authentication to be incorporated in the data itself
- ⌘ Digital signature applied to data at collection point
- ⌘ Algorithm of choice: DSA (DSS)
- ⌘ Key management centered on X.509v3 certificates
- ⌘ Conflict between distribution of CA trust functionality and single signing key



Distributed trust in certificate signature

- ⌘ X.509v3 certificates issued by a single certificate authority (CA)
- ⌘ Need to distribute trust in certificate signature (CA private key)
- ⌘ Investigated solutions: functional/role-based distribution, threshold cryptography



Investigation of threshold cryptography

- ✦ “M out of N” key components cryptographically distributed
- ✦ Components must be combined for operations (e.g. signature generation)
- ✦ Impractical for signing data at collection point, but a possible certificate-signature solution



Previous threshold implementation studied

- ✦ Identrus: Cooperative peers at high level of banking, building a trust architecture between participants and their customers
- ✦ An original Identrus participant: CertCo
 - ✦ Held threshold cryptography patents



Drawbacks to threshold cryptography

- ✧ Then-current threshold RSA had some susceptibilities (Langford '96)
- ✧ No practical threshold DSA/DSS scheme at time of system design (Desmedt '97)
- ✧ Limited solution providers
- ✧ So ruled out



Chosen solution

- ✧ Retain M-out-of-N distribution of responsibility for certificate issuance in operational distribution of roles (RAO/CAO in Baltimore terminology)
- ✧ DSA certificate signing algorithm
- ✧ Enforce >1 person access to sensitive components (key stores, etc.)



Implementation summary

- ⌘ Preliminary design, laboratory pilot, initial deployment design, initial field implementation
- ⌘ Testbed trials of lab and field systems
- ⌘ "Continuous" data: DSA signature field (40 bytes)
- ⌘ "Segmented" (message-format) data: S/MIME
- ⌘ LDAP namespace rooted on organization ID



Example: Keypair initialization

- ⌘ Keypair generation at the monitoring site:
 - ⌘ >1 onsite personnel witness, sign output, send back to data center
- ⌘ Certificate issuance
 - ⌘ >1 RAO verify received signed message from the field



Proposed command-and-control solution

- ✍ Digital signature of command/control messages
- ✍ Only holders of command-and-control certificates are empowered to issue commands
- ✍ Issues:
 - ✍ Monitoring site access to CA certificates, CRLs:
Network vs. local cache
 - ✍ Logging/auditing of control messages



Limiting aspects of system

- ✍ Certain components limited in capability or not amenable to addition of signature functionality
- ✍ Limitations on PKCS#10 certificate request generation
- ✍ Tamper-evident measures at sites but sensors could be physically moved during network or power outages (suggested adding GPS receivers but it was too costly, complex or too large for the equipment footprint)



Initial results

- ⌘ Human involvement is considerable
 - ⌘ Multiple participants, "M-out-of-N" carried into all operations where possible/applicable
- ⌘ Knowledge, skill burdens are significant
 - ⌘ PKI, trust, security, OSs etc.
- ⌘ "Ease of use," user interaction issues
- ⌘ PKI components, systems still maturing



Summary and conclusions

- ⌘ Began as a technology exercise, but deployment is heavily dependent on human factors
 - ⌘ (Trust is, after all, a human perception)
- ⌘ Trust distribution is wider than at the CA alone
 - ⌘ "Preponderance of data," number of collection sites, openness of data to scrutiny
 - ⌘ These factors were present prior to system deployment; thus, while system adds significant trust measures, it is still dependent on these other attributes
- ⌘ A qualified success, reflective of the organization
 - ⌘ Data verification systems successful
 - ⌘ Participants highly motivated, educated, skilled
 - ⌘ Like Identrus (international banking), organization has necessary resources to support the system

NIH-EDUCAUSE PKI INTEROPERABILITY PROJECT PHASE THREE PROJECT REPORT

Peter Alterman, Ph.D., Russel Weiser, Scott Rea, Deborah Blanchard ¹

Introduction

In 1998, in the Government Paperwork Elimination Act, the U.S. Federal government signaled its intention to move its transactions with citizens, businesses and other governments from paper-based applications to electronic applications accessible through the Internet. The Act required all Federal Agencies to convert paper-based transactions to electronic ones, or have plans to convert them, by October 23, 2003 and explicitly engaging the issue of electronic signatures for authentication and authorization. This first notice was followed by other statutes further empowering electronic identity management and e-business, notably the E-SIGN Act of 2000 and the E-Government Act of 2002.

In 2001, the U.S. Government identified a short list of applications to serve as leaders in implementing e-government services (the Quicksilver Project); among them were two cross-cutting services: enterprise architecture and e-authentication. This paper will focus and address the use of e-authentication. The purpose of the e-authentication program was, and continues to be, to provide electronic identity services to the 24 initial e-Gov applications and to the thousands of other government business processes that may eventually be brought on line. However, specific efforts to automate program applications were under way long before the government selected the projects that today comprise the core e-Gov activities.

The e-Authentication program initially focused on authenticating electronic identity credentials presented to the government for the purposes of authorizing citizen or business access to on-line government applications systems. A second, major category of electronic business transaction, electronically-signed electronic forms, has been acknowledged, but a concerted effort to fit it into the evolving e-authentication architecture has been on hold pending successful implementation of the first priority.

Notwithstanding the focus on access to online applications, since 2000 the Federal PKI Steering Committee has funded a project to develop models and the technology necessary to allow locally-issued digital certificates to be used to sign digital versions of government forms, and for the Federal government to be able to trust and validate those certificates. This project, [the NIH-EDUCAUSE PKI Interoperability Project](#), successfully demonstrated initial proof of concept in January, 2002 and again, using more sophisticated technology, in December, 2003. Since then, the e-authentication program and an increasing number of Federal agencies have recognized the Interoperability Project as the only successful model for implementing digitally-signed electronic forms

¹ Dr. Alterman may be contacted at altermap@mail.nih.gov; Mr. Weiser at rweiser@trustdst.com; Mr. Rea at srea@trustdst.com and Ms. Blanchard at dblanchard@trustdst.com.

processes that relies on federated identity management using X.509v3 digital certificates.

Authentication vs. Authorization

Properly managed PKI X.509v3 certificates are excellent authentication tools, especially for signing electronic documents. Given proper engineering design, they may also be useful tools to authorize access to online systems. For the purposes of Phase Three, we chose to preauthorize uploading of signed forms to the Automated Receipt Server to model secure processes found in many citizen- and business-to-government transactions. To do this, we created a small database in the Automated Receipt Server which was consulted when a signed document was presented for upload. This is a simple and straightforward method of linking authentication and authorization tools, but it has the drawback of requiring preauthorization, which can be burdensome and which adds a requirement to ensure that the same credential is always used to submit a form. A self-contained method of identifying authorization, using a business process-generated attribute, would ensure greater portability of credentials and simplify credential management at both the end user and subscriber sides. Other solutions are readily imagined.

GPEA Compliance

For the purposes of this phase of the Project, we chose to incorporate compliance with Government Paperwork Elimination Act (GPEA) requirements to demonstrate the ability of the signed forms model to satisfy statutory requirements for electronic government services. GPEA is the foundation law that requires Federal agencies to allow individuals or entities that deal with these agencies the option to submit information or transact business electronically with the agency, when practical, and to maintain records electronically, when practical.

Procedures have been identified for agencies to follow in using and accepting electronic documents and signatures, including records required to be maintained under Federal programs and information that employers are required to store and file with Federal agencies about their employees. These procedures reflect and are to be executed with due consideration of the following policies:

- a. maintaining compatibility with standards and technology for electronic signatures generally used in commerce and industry and by State governments;
- b. ensuring that electronic signatures are as reliable as appropriate for the purpose in question;
- c. maximizing the benefits and minimizing the risks and other costs;
- d. protecting the privacy of transaction partners and third parties that have information contained in the transaction;
- e. ensuring that agencies comply with their recordkeeping responsibilities for these electronic records. Electronic record keeping systems reliably preserve

the information submitted, as required by the Federal Records Act and implementing regulations; and

- f. providing, wherever appropriate, for the electronic acknowledgment of electronic filings that are successfully submitted.

GPEA defines electronic signature as a " . . . a method of signing an electronic message that:

- a. identifies and authenticates a particular person as the source of the electronic message; and
- b. indicates such person's approval of the information contained in the electronic message."²

This definition is consistent with other accepted legal definitions of signature. However, GPEA does not endorse one form of electronic signature over another, e.g., signing with a PIN versus using a digital signature. However, agencies and organizations are strongly encouraged to perform a risk assessment to determine which form of electronic signature best mitigates the risk to the agency. For the NIH-EDUCAUSE PKI Interoperability Project, the method utilized for digital signatures utilized X.509 digital certificates that were issued by the research institutions.

It is important to note the second part of the definition for an electronic signature, which requires a mutually understood, signed agreement between the person or entity submitting the electronically-signed information and the receiving Federal agency. Most often this can be accomplished by using a document referred to as a "terms and conditions" agreement. These agreements can ensure that all conditions of submission and receipt of data electronically are known and understood by the submitting parties. This is particularly the case where terms and conditions are not spelled out in agency programmatic regulations.

Products and Services

For Phase Three of the Project, the following products and services were used:

- Infomosaic SecureSign and Infomosaic SecureXML products as signing and validating tools for both the end user desktop and for the Automated Receipt Server;
- Certificate Arbitration Module (CAM) version 4.0, Release Candidate 4;
- Microsoft MSXNL 4.0 SP2 Parser and SDK;
- Microsoft IIS 6.0
- Microsoft Access 2002
- Persits Software AspEmail 4.5 Component

Participating colleges and universities used the following PKI CA products or services:

- Locally-developed CAs based on OpenSSL (two unique implementations);

² The Government Paperwork Elimination Act, section 1709(1)

- Locally-implemented iPlanet CA;
- Digital Signature Trust/Identrus-issued TrustID X.509v3 digital certificates;
- VeriSign-issued X.509v3 digital certificates from a local subordinate CA.

The Automated Receipt Server and the Federal government used ACES X.509v3 digital certificates issued by Digital Signature Trust/Identrus. The prototype Higher Education Bridge CA operated with the RSA Security Keon CA product, version 5.7p1. The prototype Federal Bridge CA contained the CA products from the following vendors: Entrust, RSA Security, BTrusted (formerly Baltimore Technologies), and Microsoft .Net CA.

Partners

In addition to EDUCAUSE, the following academic institutions are participating members of the Interoperability Project with the U.S. Federal government:

- Dartmouth College
- University of Alabama, Birmingham
- University of Wisconsin – Madison
- University of California, Office of the President
- University of Texas – Houston Health Science Center
- University of Virginia

For contact information at each of these schools, please check the Project website, <http://pki.od.nih.gov>.

Accomplishments

The accomplishments of the Interoperability Project may be divided between trust infrastructure development and PKI-enabling an electronic forms business process. During its tenure of operation, the PKI Interoperability Project has successfully created and demonstrated:

1. a certificate path discovery and validation infrastructure for assessing the legitimacy of digital certificates issued by a wide variety of PKIs and CA products;
2. an operational PKI bridge pathway between the prototype instance of the Federal Bridge CA and the prototype instance of the Higher Education Bridge CA, funded by and operated by EDUCAUSE, a nonprofit association whose mission is to advance higher education by promoting the intelligent use of information technology;
3. resolution of multiple certificate configuration and directory interoperability problems;
4. the ability for faculty and staff at academic institutions to download, fill out, sign (twice) and send XML forms to a U.S. Federal government Automated Receipt Server and to obtain an automated email acknowledgement of acceptance;
5. the ability of an Automated Receipt Server to acquire and test an XML version of a standard U.S. government form, SF-424, and to obtain an email acknowledgement of acceptance,

6. the ability of the Automated Receipt Server to automatically validate the affixed digital certificates, in the process discovering certificate validation paths to four different academic institutions using four different CA products and services through the Federal Bridge - Higher Ed Bridge pathway and to return a correct status to the server using the return path;
7. the ability of the Automated Receipt Server to send a digitally signed report containing a copy of the receipt and validation transaction, the certificates validated and a copy of the form to an audit log that satisfies the requirements of the U.S. National Archives and Records Administration for vouching for and archiving records of electronic transactions with the U.S. Federal government.

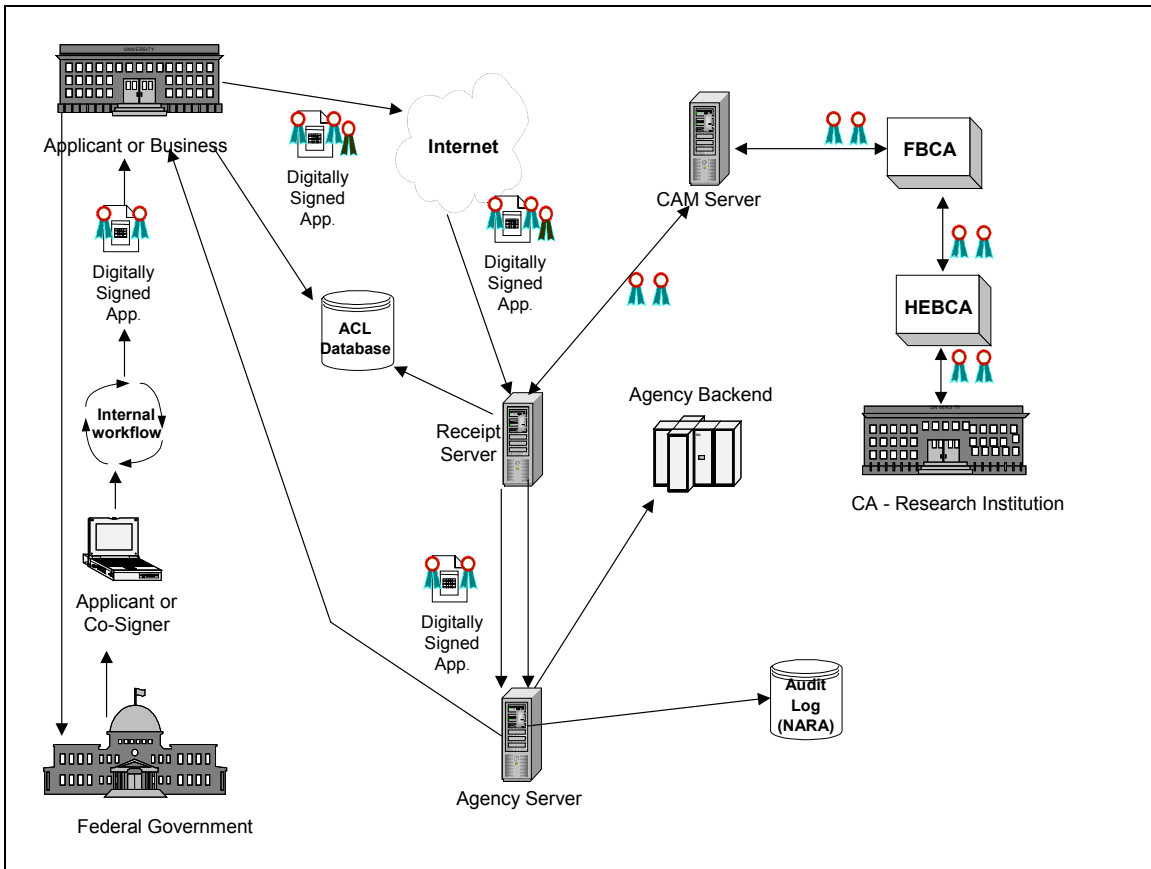
As a result of the success of the NIH-EDUCAUSE PKI Interoperability Project, a number of U.S. Federal Agencies are beginning to adopt all or part of its model to implement e-government business processes based on digitally-signed electronic forms. In addition, the Project has been recommended as a model for supporting the e-Forms initiative of the U.S. Federal government.

In the proposed Phase Four of the Interoperability Project, additional security and functionality elements will be added, especially automated parsing of the received XML form to a back-end database and encryption of all transactions within the model. Also, alternatives to the current validation tool will be evaluated.

Concept of Operations

A graphical representation of the Concept of Operations of the Interoperability Project – Phase Three appears below. Among the foundation assumptions are:

- that the document be standalone to satisfy document privacy and ownership requirements and to allow them to be managed by the end user in his or her unique environment;
- that all elements are standards-based to the extent possible at this time in the commercial environment (and that the least amount of unique code is developed and used);
- that electronic identity credentials (X.509v3 digital certificates) are issued by the institutions to which the end users belong;
- that the Federal government trust the institutional certificates at the test level of assurance;
- that, for purposes of this demonstration, one digital certificate will be preauthorized to upload the signed and completed XML form to the Automated Receipt Server, although many alternative scenarios for authorization are possible; and
- that a version of CAM 4.0 (see below for details) is the path discovery and validation tool used.



Form Conversion

To conform to emerging Federal standards for electronic forms, Phase Three replaced the Phase Two Microsoft Word template version of a PHS 398 “Application for Research Grant” with an XML version of an SF-424 “Request for Federal Funding.” The SF-424 was chosen because it was selected to be the foundation form for government-wide electronic grant applications by Grants.gov, and even though the Interoperability Project is completely separate from that project, using the same form demonstrated the broad applicability of digital certificate signing to cross-agency electronic government initiatives. The standard SF-424 has just a single signature for attestation, whereas the XML version used in this Project provides for signer and co-signer attestations. Therefore, a slight modification to the SF-424 was incorporated to demonstrate multiple signature capabilities.

The conversion of the SF-424 to *standards-based* XML was performed by mapping each requested data item in the original SF-424 form to an associated XML element and then constructing an XML schema logically patterned after the physical layout of the SF-424. The resulting schema consisted of a root element <Signed_Doc> with 3 complex-type sub-elements:

1. <TBS> (referring to the ToBeSigned portion of the document) contained all data elements being attested grouped into logical data records

2. <Signature1> the details about the first signer making attestation
3. <Signature2> the details about the co-signer making attestation

Once the schema for the Test XML SF-424 was completed, an investigation of available COTS based XML signing tools was conducted to determine the most suitable product for presentation of the Test XML form and execution of the required digital signatures. Note that the choice of signing tool was separate from the forms conversion process. Form creation and presentation/signing are standards-based, not product-specific. InfoMosaic's SecureSign Desktop was chosen (refer next section for details) for this aspect of the Project.

An InfoMosaic signature template was then designed to present the Test XML form to the user with a familiar HTML form interface. The signature template (*.tss file) is itself an XML file with a root element <SecureSignTemplate> containing minimally 6 complex type sub-elements:

1. <Header> where template name and identifying information are contained along with control elements for the form such as number of signers, hierarchy of signers and permissions for duplicate signers
2. <SchemaData> where an XML schema may be placed for validation purposes on the XML data (not used for the Project)
3. <XmlData> where the XML elements requiring data that is to be signed are defined (the entire Test XML document was encapsulated in a CDATA element here for the Project)
4. <HTMLData> where the presentation layer for the XML elements requiring data that is to be signed are defined. There is a complex type <HTMLData> element for each digital signer of the document (for the purposes of the Project there were two such elements created as individual HTML documents each encapsulated in a CDATA element). The purpose of having an <HTMLData> element for each digital signer is to allow for explicit enabling and disabling of HTML form fields in accordance with the different XML data population requirements for each signer.
5. <XPathData> where the respective elements to be signed are specified for each signer of the document (for the Project, the first signer attests the <TBS> and <Signature1> elements, while the co-signer attests <TBS>, <Signature1>, <Signature2>, and first signers signature [D-Sig]). NOTE: the co-signer is attesting the first signer's signature – this represents nested digital signatures and not just a peer signature to the original.
6. <HTMLMap> where the HTML form elements from the <HTMLData> element are mapped to XML data elements in the <XmlData> element. This allows the XML data values to be populated in the <XmlData> element from the corresponding HTML form data collected in the <HTMLData> element upon initiation of an appropriate event.

When the template is used with the InfoMosaic SecureSign application, the n-th user is presented with an HTML form created from the n-th <HTMLData> element of the template and their input is saved to the <XmlData> element based on the <HTMLMap> specified transformations upon initiation of an appropriate event (such as a Save operation). Control elements in the <Header> element determine how many signers there

are and whether file attachments are permitted. Control elements in the <XPathData> element determine whether digital signatures are nested signature or peer signatures. The SF-424 template for the Project is available from the Project.

When the InfoMosaic SecureSign application is used to provide template-based application of digital signatures, the <SignatureN> element (where N is the numeric value of the n-th signer) must be present and contain two child elements <SignatureDetailsNameN> and <SignatureDetailsDateN>. At the point of signing, the commonName attribute of the certificate to be used to verify the signature is populated in the <SignatureDetailsNameN> element and a corresponding timestamp (based on the local computer time) in the <SignatureDetailsDateN> element.

When the InfoMosaic SecureSign application is used to provide template based application of digital signatures, the template is base-64 encoded and included in the D-Sig XML output as a <SignedObject> element. Any document file(s) to be used as supporting documentation for the SF-424 submission may also be base-64 encoded and included in the D-Sig XML output as additional <SignedObject> elements.

The application of each digital signature appends a <Signature> element to the D-Sig XML output. The InfoMosaic SecureSign application can be configured to include CAM-based validation results for the signing certificate as an authenticated attribute of the signature. This is achieved by making a CAM validation request at the point of signing for the certificate chosen by the user to sign and including the response from CAM in the signature.

Signing Tool

From a number of COTS signing tools available, we selected Infomosaic SecureSign, because at the present time only SecureSign (in both desktop and server versions) implements the recently-adopted D-Sig standards for electronic document signing. In addition to being able to sign, or affix, multiple digital certificates to an XML form, or sign any file type supported by the operating system, SecureSign is able to read, display and write to templated XML forms. Thus, the same tool used to sign and validate the digital certificates is also the tool used to complete the form. This is a convenience, but does not preclude users from using other XML display applications. Information on InfoMosaic SecureSign may be found at www.infomosaic.net.

SecureSign is CAM-enabled and therefore is able to communicate with the CAM 4.0 Release Candidate 4 path discovery and path validation tool used to link certificate validation queries to the issuing PKIs through the FBCA-HEBCA mesh. (CAM is the Certificate Arbitration Module, created by the U.S. Federal ACES program; beta version 4.0 includes the Discovery and Validation Engine developed as part of Phase Two of this project by Mitretek Systems, Inc.) SecureSign was the only application discovered that provided for inclusion of certificate validation responses from CAM as an authenticated attribute of the signature. Including the CAM response as an authenticated attribute helps us to generate self-contained transactions that comply with the National Archives and Records Administration (NARA) requirements for processing and storage of electronic

records of electronic business transactions. SecureSign also currently supports native OCSP validation, SCVP validation and CRL-based validation of certificates.

The output from the InfoMosaic SecureSign Desktop application is a gzipped D-Sig XML document, making for compact storage and reduced bandwidth requirements when submitting to the Automated Receipt Server. Initially, the SecureXML server product only accepted input in the D-Sig XML format for verification and validation purposes. Due to the requirements of the project, InfoMosaic added gunzip functions to the server product so that the server could easily handle compact files submitted from the desktop product.

During the testing phase of the project, we discovered that some signed, text-based documents created for archival purposes by the SecureXML server could not be verified at a later date. Yet, when the signed document data that was being flagged as corrupt was compared to the original data used to create the document it appeared to be identical. We discovered that different character encoding was being used in creating and displaying the signed data than was used for the original data. To overcome this, InfoMosaic provided some base-64 encoding/decoding routines in their product and these were used to ensure that data remained in a consistent format by base-64 encoding the data prior to signature. This allowed for proper creation and verification of signed text-based documents.

For organizations that already have an HTML-based form and are looking to add digital signing capabilities, InfoMosaic also offers a template form designer product that will automatically create a SecureSign Template file (*.tss) for use with the Desktop product. This product was developed as a result of the template based signing capabilities incorporated in the SecureSign product for the purposes of this project.

Clearly, substantial product development by InfoMosaic was part of Phase Three. A brief summary of the modifications to SecureSign include:

1. Template based XML/HTML Form signing support in SecureSign;
2. CAM based certificate validation in SecureXML and SecureSign;
3. Development of SecureForm Designer product, a XML/HTML form designer for SecureSign;
4. Addition of gunzip feature in SecureXML;
5. Addition of various Base64 encoding/decoding APIs to SecureXML;
6. Netscape 7.X integration for signing with certificates in Netscape browsers;
7. Redesign of the SecureSign GUI making it more document oriented;
8. Master Cosigner concept - allowing a master cosigner to delete previous signatures and modify previously entered data in SecureSign;
9. Saving of unsigned documents for future completion and signing in SecureSign;
10. Signed document auto save feature in SecureSign.

Fully licensed copies of SecureSign for the desktop were provided to institutional participants. A fully licensed copy of the server version of SecureSign was employed by the Automated Receipt Server to enable automatic validation of received signed forms.

Bridges

The certificate trust infrastructure of the Project is based upon interoperation of two PKI bridges, the prototyping instance of the [Federal Bridge CA](#) and the prototype Higher Education Bridge CA. The bridge-bridge interoperability details have been published by NIST and in the report of Phase Two of the Project at the First Annual PKI R& D Workshop at <http://www.cs.dartmouth.edu/~pki02/Alterman/>.

Path Discovery and Validation Tools

Path discovery and validation through a bridge can be a complicated process – from the validation trust anchor up to the appropriate cross certificates, to another cross certificate and back down to the issuer of a the entity certificate being validated, including validating all of the certificates in the discovered path. This is further complicated by inherent issues with the bridge environments. In general, too many certificate extensions and too many certificate options make path processing too complex, too difficult, and too confusing. For example, too many options in the certificate profiles used by Bridge CAs (BCAs) introduce many complexities in the certificate path discovery and validation process. Another complication for a BCA environment is key rollover of the CA certificate and the creation and signing of CRLs. This also complicates path discovery and further complicates the validation process.

The path validation tool that was utilized in this Project was the Certificate Arbitration Module (CAM) that was enhanced with a discovery and validation engine (DAVE). CAM/DAVE was created as open source, “government-off-the-shelf” (GOTS) software. We used a beta version of this combined product known as CAM 4.0 Release Candidate 4. The package is comprised of the following public domain libraries:

- ◆ OpenSSL-0.9.7c from OpenSSL.org,
- ◆ SNACC ASN.1 Compiler created by DigitalNet,
- ◆ the Certificate Management Library (CML), version 3.2 created by DigitalNet for the Department of Defense, and
- ◆ Netscape Libraries and DLLs used for http-based or LDAP queries for AIA extensions and CDP extensions

Near the end of Phase Three, new protocols and commercial products have been announced that claim to be more robust than the CAM/DAVE validation toolset. We plan to test these in Phase Four of the Project. Additionally, updates to CAM 4.0, incorporating improvements to the CML, have recently been delivered and they, too, will be tested.

Directory

A discussion of directory issues was published in the Phase Two Report referenced above. At that time, we discussed the dependency of bridge environments on the ability to find and retrieve CA certificates, Cross Certificates, Certificate Authority Revocation Lists (CARLs), and certificate revocation lists (CRLs) to enable path discovery constructions and validation. Currently, the FBCA and HEBCA environments rely on the directory infrastructures for their proper operation. These environments differ. The FBCA Directory is based on the X.500 directory infrastructure (X.500 DSP protocol) to chain requests and knowledge reference information automatically to other external distributed X.500 directories that are participants within the FBCA environment. The requested objects are returned through the DSP protocol to the requesting directory back to the original requesting Path Discovery and Validation (PDV). The HEBCA, on the other hand, uses LDAP directories and LDAP V2 referrals to facilitate referral of the requesting client to another participant in the HEBCA environment via a URL-based process (smart referrals) to the actual institution's directory holding the needed PKI-based objects.

Both environments leverage LDAP V3 as the primary client access protocol for querying the directories, although the FBCA's X.500 directory also supports the X.500 Directory Access Protocol (DAP) clients for backward compatibility to agency applications. The HEBCA Registry of Directories (RoDs) leverages a rather simple and innovative use of an LDAP directory referral mechanism to provide centralized "smart referral" to the HEBCA participants' directories, which actually contain the PKI Objects need for PDV.

Both the FBCA and the HEBCA directories for this Project have been in use for several years (at least in the prototype BCA environment) and have had limited use. Several items of interest have come from the operation of these directories and the shifting perception of the how both models might play in BCA environments in the future. The FBCA has taken several steps to increase the FBCA flexibility. Most recently, the FBCA directory service was changed to the ISODE M-vault directory, which supports chaining LDAP referrals on behalf of PDV queries to the FBCA directory. This opens up the possibility of the FBCA directory containing smart referrals while still allowing X.500 queries to the FBCA entries with referral to be resolved on behalf of the requesting PDV.

Certificate chaining and locating objects in these directories is easier if a digital certificate utilizes the *AIA* extension. Without an *AIA* extension in the certificate, the issues related to chaining and locating objects become significant. The certificate profiles needed to support this more generic method of finding caCertificates and crossCertificatePair should include the use and population of the *AIA* extension. Very little software makes use of the *AIA* extension; however, DAVE and CAM both use the *AIA* extension if it is present. If an HTTP URL form is present, DAVE will bypass directory lookups and use HTTP directly. If an LDAP URI form is presented to DAVE, the module directly queries the given LDAP server for the given distinguished name (DN) and associated attributes and values; the same logic applies for URL-based CRL distribution point (CDP) fields to retrieve CRLs and ARLs.

An example of AIA and CDP is highlighted in the following hierarchy³.

Root CA (straight SubjectDN retrieval)

- *SIA extension* → URL=
ldap://ldap.trustdst.com/cn=DST ACES Root CA X6,ou=DST ACES,o=Digital Signature Trust,c=US?cACertificate;binary,crossCertificatePair;binary

SubCA cert

- *AIA extension* → URL=
ldap://ldap.trustdst.com/cn=DST ACES Root CA X6,ou=DST ACES,o=Digital Signature Trust,c=US?cACertificate;binary,crossCertificatePair;binary
- *CDP extension* → URL=
ldap://ldap.trustdst.com/cn=DST ACES Root CA X6,ou=DST ACES,o=Digital Signature Trust,c=US?certificateRevocationList;binary,authorityRevocationList;binary

EE cert

- *AIA extension* → URL=
ldap://ldap.trustdst.com/cn=DST ACES Unaffiliated individual CA A3,ou=DST ACES,o=Digital Signature Trust,c=US?cACertificate;binary
- *CDP extension* → URL=
ldap://ldap.trustdst.com/cn=DST ACES Unaffiliated individual CA A3,ou=DST ACES,o=Digital Signature Trust,c=US?certificateRevocationList;binary

Registry of Directories

The Registry of Directories (RoD) is a centralized list of pointers to other directories within an Internet domain. First fielded by the Internet2 Middleware Initiative for the U.S. National Science Foundation, we created a RoD for the “.gov” space and connected it to the existing Internet2 RoD used for the “.edu” space. Each RoD contains pointers to PKI directories that support the participating CAs in government and higher education. For more information on Registries of Directories, see the Phase Two Project Report referenced above.

Open Issues for the Registry of Directories

- The referral URI used in the smart referrals of the RoD must be pre-escaped. In other words, adherence to the URI definition rules must be strictly followed such that space characters must be translated to the %20 in the URI.
- Referral management will require institutional administrators to be aware of changes to the local directory tree that could affect RoD smart referrals. The LDAP Browser/Editor version 2.8.2 by Jarek Gawor was utilized for the creation of the smart referrals in the RoD. This version of the LDAP Browser/Editor was

³ Note that the AIA and CDP may simply be populated with HTTP based URLs in the case of HTTP AIA a .p7b file works for CA and cross certificate pair certificate as separate der encode binary certificates. The HTTP base CDP would have a .crl file of the CRL.

used as the native administration interface of the directory server was found to be cumbersome.

It would probably be wise to write a simple tool or script to provide a subjectDN, Institutional Directory IP address and Port. This tool could then be used to build the Smart Referral as a LDIF file that would easily be imported in the correct RoD. This would simplify the management of the RoDs and reduce errors.

Archive

One of the requirements of the Project is to archive transactions in compliance with NARA guidelines. These state, "If an electronically signed record needs to be preserved, whether for a finite period of time or permanently, then the agency needs to ensure its trustworthiness over time."⁴ Reliability, authenticity, integrity and usability are the characteristics used to describe trustworthy records from a records management perspective. Each of these characteristics was considered when implementing the archive strategy for this Project:

- ◆ Reliability – a reliable record is one in which content can be trusted as a full and accurate representation of the transactions, activities, or facts to which it attests and can be depended upon in the course of subsequent transactions or activities. To meet the goal of reliability, the Project implementation creates an XML archive record consisting of the submitted form, validation responses on each of the signing certificates, signed with the Receipt Server's own archive certificate and including a timestamp as an authenticated attribute.
- ◆ Authenticity – an authentic record is one that is proven to be what it purports to be and to have been created or sent by the person who purports to have created and sent it. To meet the goal of authenticity, the Project implementation utilizes the PKI reliance properties of the signing certificates included in the transaction as proof of origin and the PKI reliance properties of the server's own certificate as proof of acceptance into the system.
- ◆ Integrity – the integrity of a record refers to it being complete and unaltered; it is necessary to protect records against alteration without appropriate permission. To meet the goal of integrity, the Project implementation utilizes the PKI integrity properties inherent when digital signatures verify.
- Usability – a usable record is one that can be located, retrieved, presented and interpreted; any subsequent retrieval and use of the record should imply a direct connection to the business activity that created it. To meet the goal of usability, the Project implementation utilizes an appropriately indexed database to store the digitally signed archive records.

The Project utilized a Time-Contextual approach to ensure the trustworthiness of its electronically signed records over time. This was done by maintaining adequate documentation of the record's validity, such as trust verification records (signature verification and certificate validation), gathered at or near the time of record signing. This is achieved by including the CAM response in the archive documents as an authenticated

⁴ Records Management Guidance for Agencies Implementing Electronic Signature Technologies – NARA Modern Records Program, Office of Records Services, October 18, 2000

attribute of the archive XML document's signature. This approach was preferred to a Time-Independent approach since the Time-Contextual approach is less dependant on technology and much more easily maintained as technology evolves over time for records that have permanent or long-term retention requirements.

The Project implemented the following steps to ensure trustworthy records relating to electronically signed transactions:

- Created and maintained documentation of the systems used to create the records that contain the signatures
- Ensured that the records were created and maintained in a secure environment that protected them from unauthorized alteration or destruction
- Implemented standard operating procedures for the creation, use, and management of these records and maintained adequate written documentation of those procedures.
- Ensured electronically signed SF-424 trustworthiness by the following:
 - Stored the original digitally signed SF-424 form in the archive XML document
 - Digital signature on archive XML document included authenticated timestamp as part of the signature
 - Archive XML document included digital certificate for verification purposes for each signatory on the original digitally signed SF-424 form
 - Archive XML document provided for signature verification at any time for each signatory on the original digitally signed SF-424 form
 - Archive XML document included certificate validation result (from CAM) for each signatory on the original digitally signed SF-424 form and the receipt signer's own certificate validation result and an authenticated attribute of it's signature
 - Long-term integral storage of all of the above items will be achieved by optical media back-up of the archive database.

Email Receipt to the Submitter from the Server/Recipient

In support of GPEA requirements, one of the objectives of the Project was to have the SF-424 form received via the Automated Receipt Server with an automated process for verifying the signatures, validating certificates, and confirming the receipt of SF-424 to the submitter, and the signatories via an e-mail message. The following mandatory requirements for automated receipting, verification, validation, and notification were addressed:

- 24 hour real-time processing of submitted SF-424;
- SF-424 forms processed in real time as they are uploaded to the Automated Receipt Server;
- Submission received confirmation via HTML response message and email;
- Automated digital signature verification;
- Automated Certificate Validation via CAM;

The Automated Receipt Server is a web-based service that allows registered users to connect and upload their Test SF-424 XML forms. If the submitted document is not the

correct format, or if the submitter is not the co-signer on the document, then the submission file is rejected. The Automated Receipt Server uses email to send file upload receipts to each of the Test SF-424 XML Signatories – assuming that their email address is contained in the DN of their certificates; else only to the registered submitter - when processing of the submission is complete.

An administrator email account was created for the Project on the NIH mail server, allowing the Automated Receipt Server to create emails from a valid NIH email address. Access to the Administrator email account is assigned to the NIH personnel responsible for administering the Project site.

The Applicant's certificate and Co-signer's certificate are extracted from the submitted SF-424 document by the SecureXML server without any human intervention. The respective signature blocks are then verified utilizing the corresponding public keys. Once a signature is verified, the respective certificate is handed to the CAM for certificate validation. Upon successful verification of the signatures, and validation of the associated certificates, a confirmation e-mail message is sent to the certificate holders (providing there is an email attribute in the signing certificate subject DN) using the e-mail address contained within the digital certificate. If no email address is found in the certificates subject DN, a confirmation e-mail message will be sent to the registered e-mail address of the submitter stored in the access control list.

Certificate-Based Access Control

The Project used a web-based upload function to allow Project participants to submit forms. A certificate-based access control list and SSL Mutual Authentication process control access to the Project site submission service. Authorized Test SF-424 form co-signers are the only individuals who may submit forms and they are required to be pre-registered in order to do so. As a part of the registration process, the prospective submitter will be asked to nominate which certificate they will use to authenticate themselves (it must be the same certificate they use to co-sign their documents) and it is recorded in the Access Control List. The Project was originally designed for an Administrator to review all registration requests and approve or disapprove the registration, but currently operates in an auto-approve mode on all registration requests.

When a registered submitter connects to the submission service they are asked to authenticate themselves via 128-bit SSL Mutual Authentication. If the certificate chosen by the submitter for mutual authentication is not from one of the NIH trusted PKIs (FBCA cross-certified, HEBCA cross-certified, ACES or TrustID PKI), the SSL server trust list is configured for each participating Institution by installing and trusting the appropriate certificate chain for their end entity certificates, then their connection to the service is rejected. After a successful connection and a local file are submitted, the server checks the format of the file: if it is not a signed SecureSign XML document, then it is rejected. If the document is in the correct format it is saved to the server cache and the document is verified and parsed to extract the two signer's certificates. If the document does not verify or if the name of the co-signer on the document does not match the name

of the submitter, then the user is notified and the document rejected. If the co-signer is the submitter, then the two signing certificates are validated via the CAM. If either certificate fails validation, then the document is rejected. If the certificates validate, then an archive XML record is created and the transaction recorded.

Each time a new school wishes to participate in the Project, it is necessary to obtain the certificate chain for the end entity certificates that the Authorized co-signers will use to register with and subsequently sign and submit files. This is necessary because the access control mechanism for the Project site is certificate based and the web server hosting the Automated Receipt Server needs to trust the end entity chain for SSL Mutual Authentication so that the Authorized Co-signer is presented with the option to select their certificate from the list presented by the browser when they connect to the Automated Receipt Server site. If the chain is not installed and trusted on the web server, then the user will not have the option to select their certificate for authentication and will be unable to connect.

Two Digital Signatures on the Form

In a continuation from previous phases, multiple digital signatures were utilized, validated, and verified for the form submittal. However, the SF-424 in a production state only requires one signature. Since in this phase we were not building a true production system, we had the flexibility to add as many signature blocks as necessary. We chose to require two signatures as a proxy for multiple signatures, to extend the model to as great a degree of flexibility as possible.

We also incorporated limited rules around the digital signatures on the form. First, the digital certificates used for digital signing were required to be different. In other words, a participant may not use the same digital certificate to create the two digital signatures. Second, the digital certificate used for digital signing was validated and verified during submission and optionally during the digital signing ceremony. This ensured for the subscriber that the digital certificate being used was valid and that neither had it expired nor been revoked. Finally, the person submitting the SF-424 must be the second signatory and must have been registered in the ACL Database, as discussed above. Registration in the ACL Database was done using a valid digital certificate.

These design decisions were made to demonstrate that using digital certificates for digital signing, for access control, and for workflow is viable in many situations. Additionally, we were able to demonstrate that digital certificate technology can satisfy fully the business rules required for government business processes.

Next Steps

Phase Three of the Interoperability Project successfully demonstrated proof of concept in December, 2003 at the EDUCAUSE offices in Washington, D.C. Since then, several different academic participants have run the demonstration independently in a variety of venues, proving that the model is successful for school-school and agency-agency

business transactions using electronic forms signed with multiple digital certificates. Phase Four is planned to enhance the security of the model and includes designs to parse the signed form into a back end form automatically, to prove the functionality of implementing digitally signed electronic forms into a larger electronic business process scheme.

Acknowledgements

Grateful appreciation for their participation in the pilot project is acknowledged to: Clair Goldsmith, University of Texas System; Jill Gemmill, University of Alabama at Birmingham; Keith Hazelton, University of Wisconsin-Madison; Eric Norman, University of Wisconsin-Madison; Robert Brentrup, Dartmouth College; Ed Feustel, Dartmouth College; David Wasley, University of California Office of the President; Bill Weems, University of Texas – Houston Health Science Center; Barry Ribbeck, University of Texas – Houston Health Science Center; Mark Luker, EDUCAUSE; Steve Worona, EDUCAUSE; Debb Blanchard, Identrus/Digital Signature Trust; Andrew Lins, Mitretek Systems; Scott Rea, Identrus/Digital Signature Trust; Russ Weiser, Identrus/Digital Signature Trust; Manoj K. Srivastava, Infomosaic; Judy Spencer, Chair, Federal Identity Credentialing Committee and Tom Turley and Frank Newman of the NIH Office of Extramural Research.

References

1. The Government Paperwork Elimination Act, section 1709
2. “Implementation of the Government Paperwork Elimination Act”, Office of Management and Budget, <http://www.whitehouse.gov/omb/fedreg/gpea2.html>
3. Final Report: EDUCAUSE – NIH PKI Interoperability Project Project, Prepared for National Institutes of Health (NIH) Office of Extramural Research (OER), Under Contract No. GS00T99ALD0006, May, 2003
4. Memorandum to the Heads of all Departments and Agencies, “E-Authentication Guidance for Federal Agencies”, M-04-04, Joshua B. Bolton, December 16, 2003
5. PKI: Implementing and Managing E-Security, Nash, Duane, Joseph, and Brink, McGraw-Hill Publishing, 2001
6. Planning for PKI, Housley and Polk, John Wiley and Sons, 2001
7. Records Management Guidance for Agencies Implementing Electronic Signature Technologies – NARA Modern Records Program, Office of Records Services, October 18, 2000
8. Report: EDUCAUSE - NIH PKI Interoperability Pilot Project, Peter Alterman, Russel Weiser, Michael Gettes, Kenneth Stillson, Deborah Blanchard, James Fisher, Robert Brentrup, Eric Norman, 1st Annual PKI Research Workshop, <http://www.cs.dartmouth.edu/~pki02/Alterman/>.

Trusted Archiving

Santosh Chokhani & Carl Wallace
Orion Security Solutions

Abstract

Digital signatures are a powerful tool for demonstrating data integrity and performing source authentication. Timestamps are a powerful tool for confirming data existence by a particular point in time. Today, the value of digital signatures (and timestamps containing digital signatures) is limited due to a lack of tools and techniques that address the problems associated with digital signatures that accrue over time, including: expiration, revocation, cryptanalytic advances and computational advances. In this paper, we describe a system concept and protocol to achieve secure storage of data for long periods with preservation of integrity. The approach uses periodically refreshed time stamps to address these problems. The techniques can be used for a wide variety of applications, including those requiring long-term non-repudiation of digital signatures. The concept and protocol are based on minimizing trust in individual system components in order to reduce the security requirements for those components and to enhance the trust in the overall system. A proof-of-concept implementation based on the ideas and protocol described in this paper has been developed and successfully tested.

1. Introduction

One of the challenges of using digital signatures is how to prove the validity of signatures well into the future when the signer's, or a related certification authority's, credentials are no longer valid or available. Trusted archiving is a process that involves the active storage of data where evidence is periodically obtained, or generated, and stored to create an unbroken history demonstrating the integrity of data from storage time to verification time. Trusted archives are a missing piece of the PKI puzzle that are required if digital signatures are to have a durability similar to paper and ink signatures.

We have designed and developed a client-server system that addresses this problem. This paper describes our work. Section 2 contains the system concept. Section 3 provides an overview of the client-server protocol for implementing the system. Section 4 describes some security considerations. Section 5 provides a summary of the implemented system*. Section 6 describes lessons learned. Section 7 describes future plans.

* The ideas and work described in this paper (including the proof-of-concept) were funded by the United States Marine Corps.

2. Trusted Archive System Concept

A trusted archive should meet the following requirements, at a minimum:

- Provide evidence to demonstrate the integrity and, optionally, the source of data after the expiration of the cryptanalysis period for related keys and algorithms.
- Provide evidence to demonstrate the integrity and, optionally, source of data if a related certification authority (CA) is no longer operational.
- Provide active controls to protect the integrity of archived information.†

The central component of the solution is a trusted archive authority (TAA). A TAA accepts data for long-term storage and is responsible for ensuring that an evidence trail is produced and stored to enable demonstration of data integrity at any point in the future. TAAs participate in client-server transactions

† Many cryptographic mechanisms (such as digital signature or HMAC) are detection mechanisms with regard to integrity and source authentication. From a practical viewpoint, a trusted archive service needs to ensure that the archived information is protected from tampering. The mechanisms described in this paper extend the detection mechanisms and are not a substitute for secure, redundant storage, tamper protection, etc.

with entities seeking to exercise the TAA's services. TAAs use current credentials to generate signed responses as part of these transactions. Clients verify TAA signatures using a trust anchor known to the client at the time of the transaction.

Upon submission and periodically thereafter, the TAA obtains or generates a new time stamp for archived data in order to account for cryptanalytic advances against hashing or signature algorithms and to account for expiration of TSA keys. This periodic acquisition of new time stamps is referred to as "time stamp refresh" throughout the remainder of this document. The amount of trust invested in a TAA can be minimized by using the services of a trusted time stamp authority (TSA) to obtain time stamps for archived data instead of generating timestamps directly.

The client-server protocol between the client and TAA is a simple set of request/response transactions that enable submission of data to a TAA and retrieval or deletion[‡] of data from a TAA. The transactions are defined in ASN.1 and, generally, are DER encoded. Cryptographic Message Syntax (CMS), defined in [CMS], is used for all digital signatures. CMS was chosen because it is an IETF standard, many products support it, it provides flexibility to apply the cryptographic services appropriate for the application, and it provides the flexibility to include time stamps, certificates, revocation information, etc. as needed. An ASN.1 based protocol was chosen due to the requirement for an ASN.1 encoder/decoder to process most PKI artifacts. An XML submission format could be defined to provide a broad entry to a TAA. Retrieval should be sufficiently rare and in need of special purpose software, e.g. for historic algorithms, to be sustainable by a single format.

The TAA is designed to securely archive information of any type and need not have knowledge of the format of archived data. Where archived data contains digital signatures that must be verifiable in the future, collection and packaging of the items required to support signature verification are the responsibility of the archive submitter. Given the likelihood that the submitter will have performed verification, this requirement is not particularly onerous and can be easily implemented by packaging the artifacts from that verification operation, e.g. trust anchors, certificates, Certificate Revocation Lists (CRLs), Online Certificate Status Protocol (OCSP) responses, Simple Certificate Validation Protocol

(SCVP) responses, etc., with the data to archive prior to submission to the TAA. Alternatively, a TAA may provide server-side verification services to simplify and streamline the process of verifying and archiving data.

Trust anchors will come and go over the course of time but always must be obtained in a trusted manner to support certificate path validation. A TAA may archive a set of trust anchors that can be provided to retrieval clients. This capability allows the retriever to validate a digital signature without having to rely on the good intentions of the original submitter. This capability also permits a functional separation in order to enhance the trustworthiness of the archival service, i.e. one TAA can be store archived data and evidence and another TAA can store trust anchors.

In summary, the system concept consists of the following:

- TAAs use digital signatures to demonstrate the integrity and source of responses from the TAA. This is primarily a concern where responses contain trust anchors.
- TAAs periodically refresh time stamps in order to protect against advances in technology that can break hash and signature algorithms and to maintain verifiability in cases of key (or certificate) expiration.
- Archive submission clients collect and submit all information (e.g., certificates, revocation information, SCVP responses, etc.) required for long-term non-repudiation of digital signatures that cover the data submitted to a TAA.
- Archive retrieval clients verify the signatures on the TAA response and on the associated archive record to confirm the integrity of the data. The retrieval client may use trust anchors from one or more of the following sources to verify signatures contained in the evidence record or in archived data itself, if the archived data was signed:
 - Trust anchors from the signed retrieval response.
 - Trust anchors obtained independently from the same or different TAA.
 - Trust anchors known to the retrieval client or obtained via other out-of-band means.

[‡] Where a TAA maintains archived data on write-only media, deletion may simply be cessation of refresh operations rather than actual deletion. Deletion is not addressed in detail in this document.

3. Trusted Archive Protocol (TAP)

3.1 Assumptions and Background

The Trusted Archive Protocol (TAP) was developed and submitted to the IETF as an Internet Draft (I-D) of the PKIX working group. The TAP I-D was a contributing factor in the formation of the IETF Long Term Archive and Notary (LTANS) working group (WG) and has served as input to the protocol being produced by that group. Activities of the LTANS WG and its relationship to this work are described in Further Research section.

TAP was designed using the following principles:

- The CMS will be used in all cases where digital signatures are applied.
- A TAA shall provide an archive submitter a response that includes a time stamp token, identifying information and a TAA-generated digital signature. Clients can verify the timestamp token to confirm the correct data was received and (presumably) archived by the TAA.
- The TAA shall verify the time stamp token received from the TSA in accordance with RFC 3161 [TSP].
- The TAA shall periodically refresh the time stamp token.
- The TAA shall provide all timestamps obtained for the archived data in the response.

The rest of this section provides a summary of the protocol defined in [TAP].

3.2 Definitions

During the development of TAP, the need arose for a common vocabulary describing the various processes and artifacts involved in archiving. The following terms were defined to meet this need:

Archived data: archived data is the data presented to the TAA by the submitter.

Archive token: an archive token is an object generated by the TAA when data is submitted and accepted for archiving. The archive token is returned to the submitter and may be used to request retrieval or deletion of the archived data and associated cryptographic information. For purposes of future retrieval or deletion, applications may treat the archive token as an opaque blob. The archive token

includes: submitter DN, timestamp token, TAA date and time upon submission and, optionally, tracking information.

Archive record: an archive record contains the cryptographic refresh history compiled by the TAA. The initial archive record is the timestamp token obtained for the submitted data. The timestamp token format is defined in [TSP] and consists of a ContentInfo object containing a TSTInfo object. Upon each refresh, the most recent archive record becomes the prevArchRecord field of a new TimeStampedData object, a timestamp is obtained for the TimeStampedData object and is placed in the timestamp field of a new ArchiveRecordData and the entire ArchiveRecordData structure placed in a ContentInfo object. The ContentInfo object serves as the new archive record. When verifying an archive record, verification terminates when the original timestamp token is verified against the archived data.

Archive package: an archive package is an object containing, minimally, the archive token, archive record and archived data. The archive package may include additional cryptographic information. Archive packages are returned during retrieval.

Figure 1 illustrates the communication protocol among the TAA and the clients.

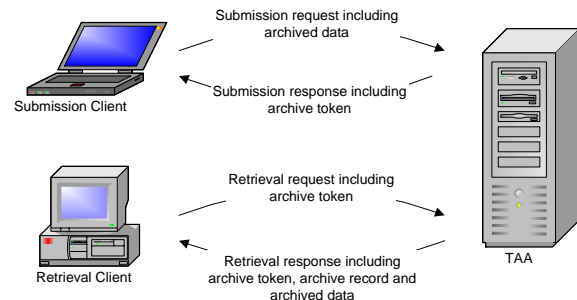


Figure 1 Client interactions with TAA

Figure 2 illustrates the archive record that is maintained by the TAA. The archive record contains nested timestamps with a timestamp covering the archived data at its innermost layer.

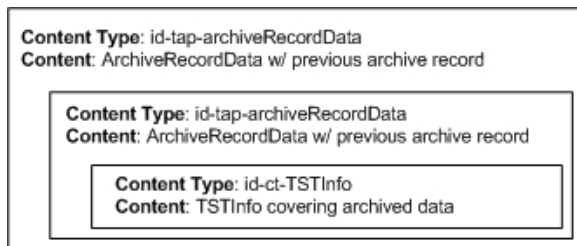


Figure 2 Archive record after two refresh operations

3.3 Protocol Summary

The [TAP] protocol defines 3 request types: submission, retrieval and deletion. The steps involved in a submission request are as follows:

- 1) A client prepares a data object for submission to a TAA. If the data object is signed, the client verifies the object and includes the necessary material to verify the object (except the trust anchor) in the object itself, e.g. in a certificate or CRL bag. Optionally, the client signs the request. The client sends the request to the TAA.
- 2) The TAA receives the request and verifies the signature on the request, if present. The TAA unpacks the data object and prepares a TSP request. Optionally, the TAA signs the TSP request. The TAA then sends the TSP request to a TSA.
- 3) The TSA receives the response and verifies the signature on the request, if present. The TSA then generates a signed timestamp token and returns it to the TAA.
- 4) The TAA stores the archived data and the timestamp token and starts the refresh clock for the archived data. The timestamp token is packaged in an archive token along with additional information. The archive token is included in a signed response and returned to the client.
- 5) The client verifies the signature on the response. The client verifies the archive token to ensure the correct data was archived by the TAA. The client may store the archive data along with the original data item, e.g. as an unsigned attribute.

The steps involved in a retrieval request are as follows:

- 1) A client prepares a retrieval request containing the archive token of the data item for which an archive record is required. The

client signs the request and sends it to the TAA.

- 2) The TAA verifies the signature on the request and confirms the requestor has access to the requested data item. The TAA prepares an ArchivePackage containing the refresh history compiled for the requested data item, packages it in a signed response and returns it to the client.
- 3) The client verifies the signature on the response then verifies the archive package. The outermost layer in the archive record is verified using a current trust anchor. Interior layers are verified to a trust anchor provided by the TAA in the archive package.

3.4 Protocol Data Formats

The section describes some of the key data formats defined in [TAP].

Archive Submission

Archive submission requests are defined as follows:

```
ArchiveSubmissionReq ::= SEQUENCE
{
    version          TAPVersion DEFAULT v1,
    submitterName   GeneralName,
    policy          OBJECT IDENTIFIER
                  OPTIONAL,
    archiveControls [0] ArchiveControls
                  OPTIONAL,
    archivedData    ArchivedData
}
```

Archive Data

Archived data, i.e. data submitted to a TAA for preservation, has the following format.

```
ArchivedData ::= SEQUENCE
{
    type    ArchivedDataType OPTIONAL,
    data    OCTET STRING
}
ArchivedDataType ::= CHOICE
{
    oid        OBJECT IDENTIFIER,
    mimeType   UTF8String
}
```

Archive Submission Response

Archive submission responses are defined as follows:

```
ArchiveSubOrDelResp ::= SEQUENCE
{
    version          TAPVersion DEFAULT v1,
    status          ArchiveStatus,
    archiveToken    ArchiveToken
}
```

```

        archiveControls  [0] ArchiveControls
        OPTIONAL
    }

```

Archive Token

Archive tokens have the following format.

```

ArchiveToken ::= ContentInfo
-- content type: id-tap-archiveToken
-- content: ArchiveTokenData
ArchiveTokenData ::= SEQUENCE
{
    submitterName      GeneralName,
    timestamp          TimeStampToken,
    curTime            GeneralizedTime,
    trackingInfo       TrackingInfos
                    OPTIONAL
}

```

The archiveControls field can be used to return information associated with a control included in the request, for example, the outcome of server-side validation or a nonce from the request. TAAs must not include controls in a response that are not associated with controls in a request. Submission clients should be able to process controls in accordance with the control definition.

Archive Record

The archive record contains a nested structure with the complete refresh history for the archived data. TAAs should store all cryptographic information necessary to verify each layer of the archive record in the certificates, CRLs and unsignedAttrs fields of the timestamp token, i.e. each timestamp token in the history should be self-contained for validation purposes under protection of the next layer in the archive record. A CryptoInfos unsignedAttrs field may be used to convey OCSP responses and/or trust anchor information. Archive record has the following format:

```

ArchiveRecord ::= ContentInfo
-- content type: id-tap-archiveRecordData
-- content: ArchiveRecordData

ArchiveRecordData ::= SEQUENCE
{
    timestampedData    TimeStampedData,
    timestamp          TimeStampToken
}
TimeStampedData ::= SEQUENCE
{
    prevArchRecord    ContentInfo,
    messageImprint    MessageImprint
}

```

The cryptoInfos field contains additional information that may be useful when verifying the archived data.

Archive Package

Archive packages are defined as follows:

```

ArchivePackage ::= SEQUENCE
{
    archiveToken      ArchiveToken,
    packageData      [0] ArchivePackageData
                    OPTIONAL,
    pollReference     [1] OCTET STRING
                    OPTIONAL
}
ArchivePackageData ::= SEQUENCE
{
    digestAlgs        DigestAlgorithmIdentifiers,
    policy            OBJECT IDENTIFIER
                    OPTIONAL,
    archRecord        ArchiveRecord,
    cryptoInfos [0] CryptoInfos
                    OPTIONAL,
    archivedData      ArchivedData
}

```

4. Security Considerations

This section provides an overview of a security analysis of the protocol.

Trust Anchors for Timestamp and Other Signature Verification on Archive Retrieval

TAAs can provide all or some of the trust anchors upon retrieval. These may include all the trust anchors required to verify the various timestamps in the archive record and/or all the trust anchors known to the TAA at the time of the archive submission (i.e., the timestamp on the archived data). The latter set of trust anchors may be useful in digital signature verification on the archived data, if the data was signed.

Trust anchors provided by the TAA upon archive retrieval are transmitted securely since they are included in the signed envelope of the retrieval response. The relying party (i.e., the retrieval client) must use a trust anchor it trusts independent of the trust anchors provided by the TAA to verify the TAA signature on the retrieval response.

The relying party (i.e., the retrieval client) can trust the TAA provided trust anchors or can ignore them. In the latter case, only the TSA (and not the TAA) needs to be trusted for the integrity of the archived data. In other words, the relying party will be able to detect the modifications made to the archived data by the TAA. Refreshing the timestamp on the archived data before the latest (i.e., most current or outermost) timestamp expires ensures this.

Algorithm and Technology Advances

In order to protect against algorithm (i.e., hashing and digital signature) compromise and/or computing technology advances, timestamps are periodically refreshed. For each timestamp token refresh, the

archived data is hashed using a current, secure hashing algorithm and a timestamp token generated using a current, secure digital signature algorithm.

Security of TSA

TAA's must be able to obtain a trusted timestamp (either by implementing timestamp functionality or by access to a timestamp service). Timestamp-related security considerations apply (see [TSP]).

ArchiveControls

ArchiveControls are optional request components that request server-side processing in addition to archiving, i.e. collection of certificates and CRLs. ArchiveControls that request alteration of the submitted data should define a response such that the timestamp contained in the archive token can be verified.

5. System Description

A trusted archive system using the requirements, concepts and protocol presented here has been developed to successfully demonstrate the concepts presented in this paper.

The components of the system are as follows

- A [TSP]-compliant Time Stamping Authority (TSA)
- A [TAP]-compliant TAA
- [TAP]-compliant TAA clients

The following diagram depicts the overall architecture of the implemented system.

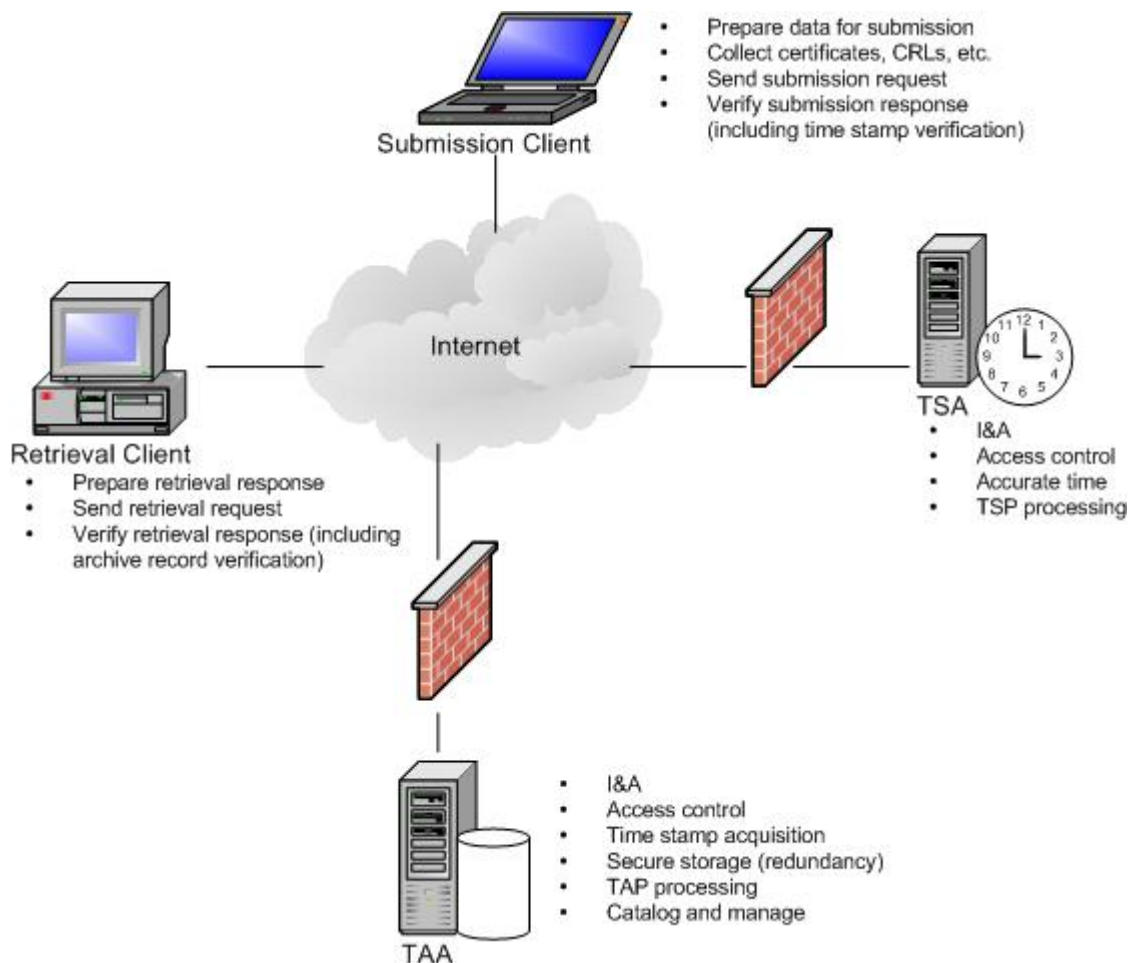


Figure 3 System Architecture

5.1 Time Stamp Authority (TSA)

The TSA is RFC 3161[TSP] compliant and is hosted on a PC running Windows 2000 Server. The TSA uses the National Institute of Standards and Technology (NIST) Internet based Network Time Protocol (NTP) service to set the Windows 2000 TSA Server system clock. The TSA interacts with TSA clients using HTTP. The TSA has a public key certificate issued by a PKI recognized by the TAA and, optionally, TAA clients.

5.2 Trusted Archive Authority (TAA)

The TAA is TAP I-D compliant and is hosted on a PC running Windows 2000 Server. The TAA has a public key certificate issued by the PKI recognized by the TAA clients. The TAA interacts with TAA clients using HTTP.

The TAA obtains the initial time stamp from the TSA upon submission of archived data. The TAA periodically refreshes the time stamps in accordance with the system concept and the TAP protocol.

The TAA catalogs archive data using the following attributes: submitter DN, time stamp token, submission date and time.

5.3 TAA Submission Client

The submission client operates on Windows workstations. The client validates the TAA signature in accordance with TAP and the time stamp token contained in the archive token in accordance with RFC 3161[TSP].

5.4 TAA Retrieval Client

The retrieval client operates on Windows workstations. The client provides an archive token to the TAA in order to retrieve the archived data. The client validates the TAA signature in accordance with TAP and the timestamp tokens contained in the archive record in accordance with TAP and RFC 3161[TSP].

A retrieval client unwinds the nested CMS package consisting of multiple nested time stamps. The retrieval client verifies the various time stamps as the CMS package is unwound using the time from the adjacent outer layer as the time of verification. The outermost layer is verified using the current time. The client determines the unwinding is complete when the innermost TSTInfo is reached. The

innermost TSTInfo is used to verify the archived data.

The retrieval client may use trust anchors provided by the TAA during archive record verification or trust anchors available locally.

5.5 Operational Considerations

The system described in this section is a proof of concept implementation. If this were an operational system additional security measures are recommended akin to the operations of a CA. It is desirable that the servers and workstation use FIPS 140-2 validated hardware cryptographic modules and Common Criteria validated operating systems and application software. The operational systems and services should use Physical, Procedural, and Personnel (P³) security controls commensurate with the security needs and perceived risks.

In addition to the computer security controls described for the TSA and TAA in System Description, appropriate boundary control product (e.g., validated to conform to [FWPP]) should be used to protect the TSA and TAA.

To enhance the security of the trusted archive service using the principle of separation of duties, consideration should be given where one TAA archives the data while another TAA trust anchors relevant to the verification of signatures on the archived data. Timestamps could be obtained from multiple TSAs to limit the damage resulting from TSA compromise. Redundant storage mechanisms should be employed to ensure that no archive data is lost due to device failure or catastrophe.

6. Lessons Learned

6.1 Metadata

One significant piece of information not included in the TAP protocol was filename and format of the archived data. This information is essential when working with material retrieved from an archive. Future versions of the protocol will include means of including a variety of metadata with an archive submission.

Metadata may also be useful in aggregating archived data over time. For example, to associate a refutation of a document with the original archived document or to associate data related by context, such as a various pieces of data in a criminal file.

6.2 Timestamp reliance

The archive record structure defined in TAP relies heavily on timestamps as defined in TSP. This has a number of potentially undesirable properties including:

- A new digital signature as part of the preservation of integrity of a digital signature
- A great degree of trust is invested in the TSA
- A one-to-one ratio of timestamps to documents to archived data objects makes development of high-performance applications difficult

Alternative timestamp structures have been defined that address these concerns by relying on the security of hash algorithms and the availability of published information, see [HOWTO] and [EFF].

6.3 Search features are important

TAP featured limited means for searching an archive. The retrieval interface was highly driven by hashes of archived data. While this works well if the data is stored with its archive token, such storage may not be the norm. Without the archive token, the effectiveness of a search is highly correlated with the submission volume of the original submitter. Search features should include means of searching based on content, metadata and/or keywords.

6.4 Auto-deletion

[TAP] defined no means for clients to define the period of time a TAA should preserve a data object. This leaves the burden on the submitter to stop the refresh process at some point in time. While this could be negotiated using the policy field, a better solution would be to provide a means for specifying the archivation period at submission time.

This leads to a need to manage the archivation period (or meta-data) post-submission. A better approach to the protocol may have been to specify a submission request, a management request and a single response type. The management request would be used to retrieve and delete archived data as well as to update meta-data, archivation period, etc. A single response format would simplify the handling of errors that are not request-specific ([TAP] defined two response types).

7. Future Directions

The LTANS WG has become quite active and is developing three standards in the area of trusted archive:

- Trusted Archive Requirements
- Evidence Record Syntax (based on [ATS])
- Trusted Archive Protocol (based on [TAP])

Another area of research is authentication and authorization for deletion and retrieval. The challenge of authentication and authorization validation in support of long-term non-repudiation can be summarized as follows:

- The identity of authorized parties may change over time. Generally, [TAP] was intended to support claims against data that occur within the memory of a person or institution where retrieval would be performed by the original submitter or by an authorized agent of an organization.
- The definition of authorization attributes may change over time and the naming of attributes may not prove to be unique in contexts that expand over time.
- The authorities such as CA or Attribute Authority (AA) may be no longer in existence.

Data formats, or data format migration, are another area of concern for long term archives. The formats of signed documents today may not be readily usable after a number of years.

Providing confirmation that a specific person generated a data item after a very long period of time is a very difficult problem. Over great periods of time it would be difficult to state with confidence that a particular signature was generated by a particular person. One approach may be to archive the information collected by a CA/TA to establish the binding between a person and a key. The need for this sort of demonstration may be very small. Notarization may be of assistance in this area. Rather than maintain evidence to demonstrate the binding of keys to members of the general population, it may be necessary to simply maintain evidence that binds keys to notaries, who generate attestations at a time when sufficient information is available to confirm the binding of a person to a key and a key to a signature. Biometric information provides another alternative for establishing a link between a specific individual and an archive record and/or an individual's key. This is an area that requires further consideration.

Other mechanisms for providing TAA functionality such as n of m splitting based on Shamir technique [SHA] that would provide a high degree of availability and integrity.

While these problems exist today in general, they are more likely to be encountered when one looks at 20 to 50 years and beyond. Solutions to these issues are a fertile area of research.

References

[ACTS] NIST Automated Computer Time Service (ACTS), <http://tf.nist.gov/service/acts.htm>

[ATS] Brandner, R., Gondrom, T., Pordesch, U. and M. Tieleman, "Archive Time-Stamps Syntax", draft-brandner-et-al-ats-00.txt, July 2003.

[CMS] Housley, R., "Cryptographic Message Syntax", RFC 3369, August 2002.

[EFF] Bayer, D., Haber, S. and W.S. Stornetta, "Improving the Efficiency and Reliability of Digital Time-Stamping", *Sequences II: Methods in Communication, and Computer Science*, pp.329-334, March 1992.

[FWPP] U.S. Government Firewall Protection Profile for Medium Robustness Environments, Version 1.0, October 28, 2003.

[HOWTO] Haber, S. and W. S. Stornetta, "How to Time-Stamp a Digital Document", *Journal of Cryptology*, Vol. 3, No. 2, pp. 99-111, 1991.

[LTES] Pinkas, D., Ross, J., and N. Pope, "Electronic Signature Formats for long term electronic signatures", RFC 3126, September 2001.

[NTP] Network Time Protocol Specification, Implementation and Analysis, Internet RFC 1305, March 1992.

[OCSP] Myers, M., Ankney, R., Malpani, A., Galperin, S. and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 2560, June 1999.

[SCVP] Malpani, A., Housley, R. and T. Freeman, "Simple Certificate Validation Protocol (SCVP)", draft-ietf-pkix-scvp-13.txt, October 2003.

[SHA] Shamir, A., "How to Share a Secret", *Communications of the ACM* 24 (1979), n. 11, (1979), 612-613.

[TAP] Chokhani, S. and C. Wallace, "Trusted Archive Protocol", draft-ietf-pkix-tap-00.txt, February 2003.

[TSP] Adams, C., Cain, P., Pinkas, D. and R. Zuccherato, "Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)", RFC 3161, August 2001.



SECURITY SOLUTIONS

Trusted Archiving

3rd PKI Research Workshop

April 13, 2004

Santosh Chokhani - chokhani@orionsec.com
Carl Wallace - cwallace@orionsec.com



Special thanks to U.S. Marine Corps Systems Command

Lt. Col. Beck

Maj. Diersen

Capt. Fillmore

Mike Henry

Norm Cobb

Agenda

- **Focus of proof of concept**
- **Goals of trusted archiving**
- **System concept**
- **Trusted archive protocol**
- **Archive records**
- **System considerations**
- **Alternative solutions and research issues**
- **Future work**

Focus of proof of concept

- **Primary target was digitally signed data**
 - **Extend the period in which digital signatures can serve as tools for demonstration of data integrity and authentication of data source**
 - Non-repudiation
- **Primary goal was development of tools to enhance current PK-enabled products**
 - **Solution defined in terms of existing standards (e.g. X.509, RFC3161, CMS)**

Goals of trusted archiving

- **Provide capability to prove integrity of data over long period of time**
 - Account for expiration of relevant keys (e.g. trust anchor, TSA)
 - Account for expiration of cryptographic mechanisms (e.g. hash algorithms)
- **Optionally, provide capability to prove source of data over long period of time**
 - Optional because not all data includes material that can be used to demonstrate the source of the data
- **Ensure archived data is not modified**
 - Cryptographic mechanisms only prove integrity; they do not protect integrity

Trusted archive system concept

- **3 principals: Trusted Archive Authority (TAA), submission client, retrieval client**
- **TAA obtains timestamps from an independent Timestamp Authority (TSA)**
 - Reduces the trust required in the TAA
 - TAA signature is relevant only when the TAA serves as a trust anchor source for archive record time stamp verification
- **4 primary data artifacts**
 - **Archived data:** data submitted to a TAA for preservation
 - **Archive tokens:** returned and verified upon submission; used to retrieve archive packages
 - **Archive record:** timestamp refresh history for an archived data object
 - **Archive package:** contains archive token, archive record, archived data, and information necessary to verify the archive record

Trusted archive system concept (continued)

- **Refreshed Time Stamp Approach**
 - **Prior to expiration of a time stamp, document is time stamped again resulting in nested timestamps**
 - Expiration means expiration of TSA public key certificate
 - **Time stamp can be refreshed for other reasons**
 - Weakness in cryptographic algorithms, advances in computation
 - **All data is hashed upon refresh using current hash algorithm to protect against cryptographic and computational advances**
 - **For each layer in refresh history, all material necessary to verify the adjacent inner layer is collected and stored in the record**
 - **Outermost layer may be verified using current trust anchors known to retrieval client**

Archive record with refreshed timestamps

Content Type: id-tap-archiveRecordData

Content: ArchiveRecordData w/ previous archive record

Content Type: id-tap-archiveRecordData

Content: ArchiveRecordData w/ previous archive record

Content Type: id-ct-TSTInfo

Content: TSTInfo covering archived data

- Refresh history for an archived data object
- Outer layers are archive record structures
- Innermost layer is an RFC3161 timestamp covering archived data object

Trusted archive protocol: submission

- 4) Client processing**
- Verify TAA signature on response
 - Verify timestamp from archive token
 - Store archive token for future use



1) Submission request

- submitter's name
- archived data
- policy (optional)
- archive controls (optional)

3) Submission response

- status
- archive token
- archive controls (optional)

2) TAA Processing

- Check authentication and authorization (optional)
- Process archive controls, if present
- Obtain (or generate) a timestamp for archived data
- Create archive token and archive record
- Store archive data and archive record
- Generate response containing archive token and archive control responses
- Sign and send response

Trusted archive protocol: archive tokens

Submitter	Timestamp Token	TAA Date and Time
-----------	--------------------	----------------------

- **Returned in submission response**
- **Includes timestamp for archive data that is signed by TSA and can be verified by submitter**
- **Uses to initiate retrieval operations**

Trusted archive protocol: retrieval

- 4) Client processing**
- Verify TAA signature on response
 - Verify archive record (including all timestamps)



1) Retrieval request

- requestor's name
- archive token (or info to initiate a search)
- archive controls (optional)

3) Retrieval response

- status
- archive package
- archive controls (optional)

2) TAA Processing

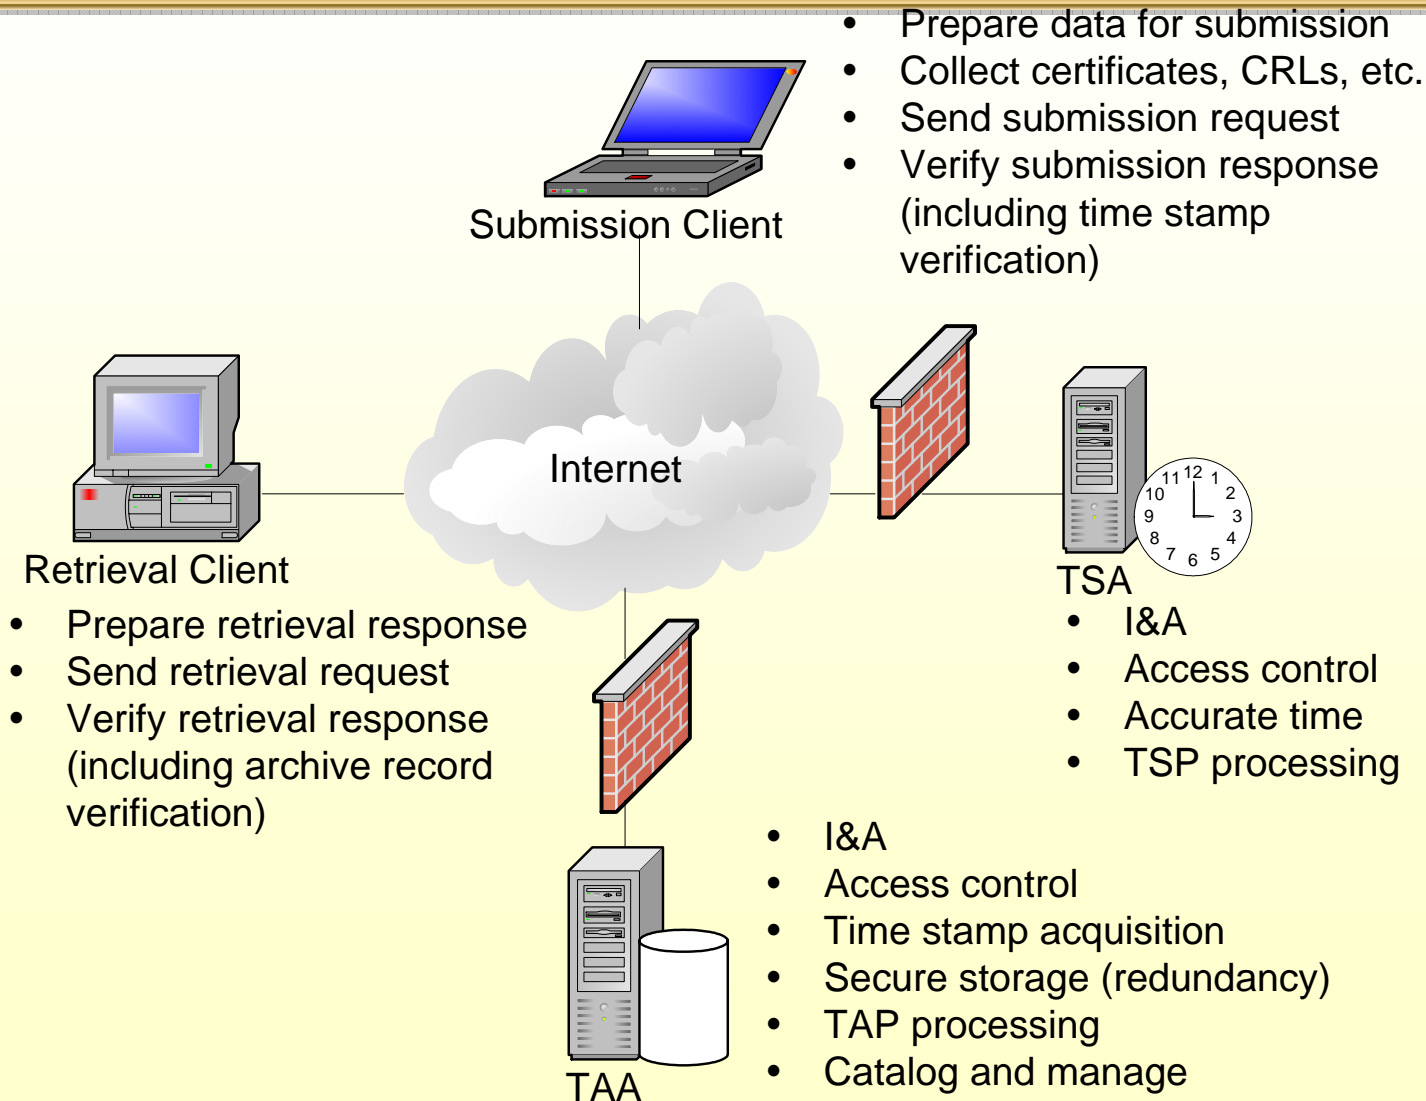
- Check authentication and authorization (optional)
- Process archive controls, if present
- Retrieve archive record and archived data
- Create archive package containing archived data, archive token and archive record
- Generate response containing archive package and archive control responses
- Sign and send response

Trusted archive protocol: archive packages



- **ArchiveRecord may include CryptoInfos, which may include trust anchors, certificates, CRLs, OCSP responses, DVCS responses and/or SCVP responses**
- **Archive record can be independent of archived data**

Trusted archive system implementation



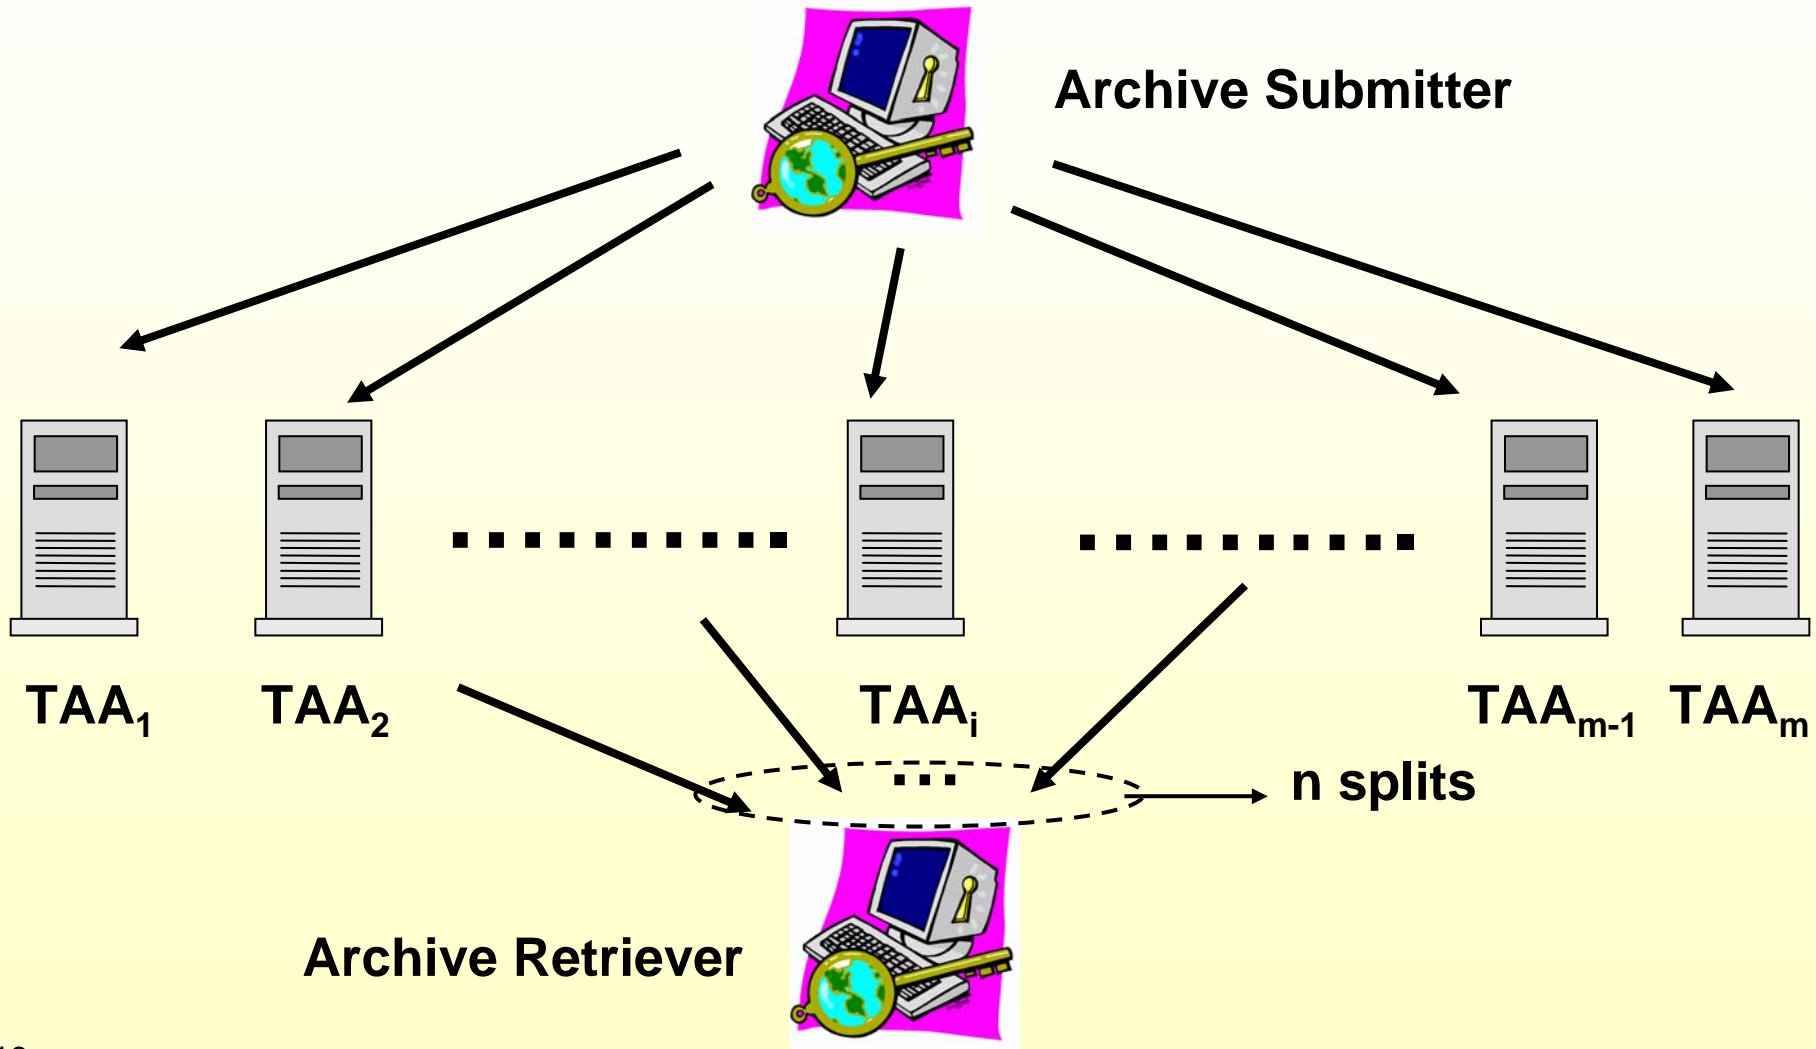
Trusted archive system considerations

- **Archived data can be unsigned, signed or time stamped**
 - If signed or time stamped, submitter should include all relevant crypto material
 - Trust anchors, certificates, CRLs, OCSP and SCVP responses, etc.
- **Confidentiality of Data from TAA**
- **TAA Threats**
 - Insider threat, Hacking, etc.
 - Distinct trusted authorities
 - TSA, TAA for data, and TAA for trust anchors

Alternative Solutions: data splitting

- **Archive Data Splitting Approach**
 - **Archive Data submitter splits the data in m shares of which n shares are required to reconstitute the data**
 - Confidentiality without key management
 - **Submitter provides 1 share to each of m TAAs**
 - **Retriever must obtain at least n shares from n TAAs**
 - **Use TLS to provide integrity (due to patent concerns about split verification) and confidentiality (to protect against eavesdropper collecting splits) for submission**
 - **n of m splits can be reconfigured based on research**

Alternative Solutions: data splitting (continued)



Alternative Solutions: data splitting (continued)

- **Issues and challenges for data splitting approach**
 - **Requires stronger I&A**
 - A rogue TAA may get $n-1$ splits and manipulate its split
 - Marketplace convergence could place n shares under single party control
 - **Metadata challenge is greater**
 - **Storage requirements and bandwidth requirements are greater (at least $m * \text{archived data size}$)**
 - **Lack of support for the approach in LTANS community**

Research issues


- **Format migration**
 - Changes in software technology may render the data formats of archived data objects obsolete
 - Non-reversible format translation will break digital signatures
 - May indicate sanctity of original data integrity is less important than initially thought and increase demand for notarization solution
- **Long-term identification of principals (e.g. authorized requesters and data source)**
 - Parties authorized to access (or manage) an archive record will change over time
 - Difficult to demonstrate binding of data source to key over long periods
 - Biometrics, PKI Registration
- **Notarization**
 - Reduces number of principals requiring long-term identification


- **Internet Engineering Task Force (IETF) Working Group (WG) established to produce trusted archive related standards**
- **Mailing-list address: ietf-ltans@imc.org**
- **Web site: <http://www.imc.org/ietf-ltans>**
- **LTANS WG Work Items:**
 - **Trusted Archive Requirements -- Internet Draft (ID) Published**
 - **Evidence Record Structure -- ID Published**
 - **Trusted Archive Protocol -- ID based on PKIX TAP ID being developed**



QUESTIONS?

List of Acronyms

CA	Certification Authority
CRL	Certificate Revocation List
I-D	Internet Draft
IETF	Internet Engineering Task Force
LTANS	Long Term Archive and Notary Services
OCSP	Online Certificate Status Protocol
PKI	Public Key Infrastructure
PKIX	IETF PKI WG
RFC	Request for Comment
SCVP	Simple Certificate Validation Protocol
TAA	Trusted Archive Authority
TAP	Trusted Archive Protocol
TSA	Time Stamp Authority
TSP	Time Stamp Protocol
US	United States
WG	Working Group


 E-filing using digitally signed Acrobat forms
or
The PKI / CA roller coaster – from hope to hype to hysteria to, finally, happiness
Ron Usher
C.E.O. - Juricert™
Staff Lawyer - The Law Society of British Columbia
For the
NIST 3rd Annual PKI R&D Workshop
Gaithersburg, MD 13 April 2004

 PKI – and I mean really big “T”



 What we often really need...






What we needed was “PKE”

- **Public Key “Enough”**
- **The right amount of infrastructure appropriate to the**
 - Risk
 - Value
 - Cost
 - Participants
 - System
- **Usually what we really need is Public Key Cryptography, with only some of the infrastructure.**




Juricert™


- Incorporated as a Canadian, for profit Federal corporation.
- Founded by the Law Society of British Columbia
- An initiative of the 14 Law Societies of Canada
- Supported by the the Federation of Law Societies of Canada (www.flsc.ca)
- Share ownership restricted to non-profit professional regulatory bodies



A credential, not certificate authority


- **Juricert is technology neutral**
 - Not dependent on a single certificate technology
 - Can be an RA, may develop a CA if needed (but very unlikely)
 - Partners with existing certificate issuing authorities
 - Can support a variety of local and national initiatives
 - Supports security system other than PKI, as appropriate for the application in question





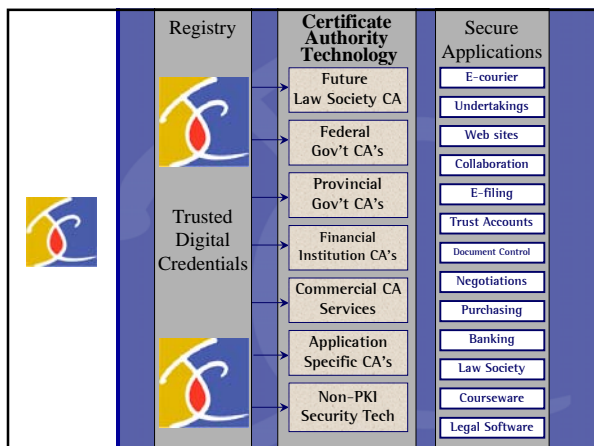
Trusted Digital Credential™

- Is a uniform digital record
 - created by the Juricert secure identification and authentication process,
 - issued by a professional regulatory body (Law Society),
 - linking a person and their professional status to an electronic identity.
- This record is used as the basis for the issuance of PKI certificates and accreditation to other internet security systems.
- Use and maintenance of this record is legally and technically controlled by the issuing professional regulatory body (such as a Law Society) through Juricert.



The Juricert I&A Process

- Registrant completes web form
- Data submitted via SSL
- Printed form witnessed, then faxed and mailed
- Fax image and data matched
- Workflow system presents image and data to PRB credential department (for professional applicants)
- Data and signatures validated against existing extensive paper and digital records
- Acceptance by PRB creates "TDC" in Juricert data base
- Ongoing moves and changes validated by PRB
- Only necessary info. released to applications



Paper forgeries are a problem

RELEASE OF MORTGAGE
BP26604
and BP132972

TERMS: Part 2 of this instr...

Belmont Ave. Fraud March 2004

Belmont Ave. Transfer

4. TRANSFEROR(S):*
ROSY PALACE INVESTMENTS


5. FREEHOLD ESTATE TRANSFERRED:*

Transferor(s) Signature(s)
Rosy Palace Investments by its
Authorized Signatory
Susan Chen
Susan Chen



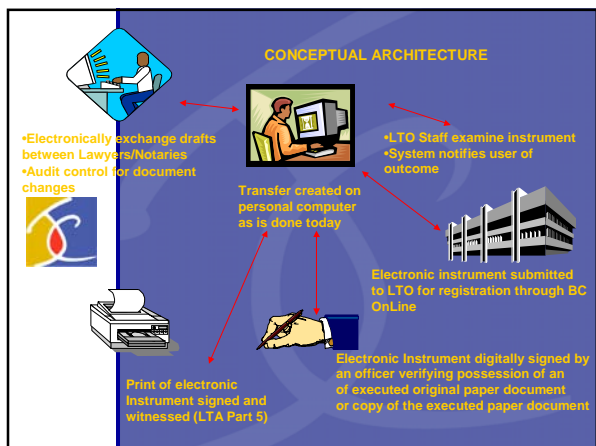
A few terms & acronyms

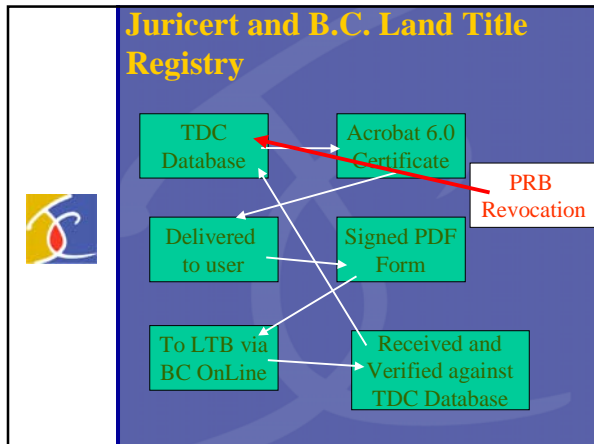
- “EFS” – Electronic Filing System
- “LTB” – Land Title Branch
- “Torrens” – the kind of legal interest in land registration system in B.C.
- “LSBC” – The Law Society of British Columbia – the statutory regulatory body for the legal profession in B.C.
- “B.C.” – British Columbia
- PRB – Professional Regulatory Body



The EFS project so far...

- First discussions in June of 1998
- EFS Committee meeting since then
- Legislation in 1999 (Bill 93)
- 3 contractors....system in “user acceptance” testing in the fall of 2003
- Bill 90 in Nov. 2003
- December 2003 LSBC/Juricert recognized as a “CA” under the 1999 legislation
- “Production Pilot” filings started 7 Jan 2004
- Continuing Legal Education programs for 1300 registrants in March, 2004
- Full Availability as of April 1, 2004





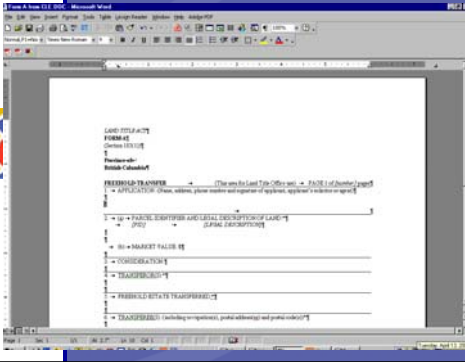
- ### Key principles implemented
- No change to property law
 - No significant change to the process and procedures of conveyancing
 - Current roles and responsibilities kept in place
 - Minimal investment by law firms – “COTS” software
 - Public interest protected by enhancing the Torrens system, keeping lawyers involved, and implementing appropriate security measures

- ### The first filing(s) – January 7, 2004
- CA2 and CA3 – test filings
 - Decided to do a line of credit mortgage for the Bank of Montreal...
 - Documents had already been prepared, so the mortgage was redone using the Acrobat Form B...
-

Roz does up the new form in Acrobat v. 6



For most firms they typically use Word to create forms



Really, it is very easy....



The mortgage form...

FORM 8_V1

[PID] **Use Schedule** [legal description]

performance of its obligations in accordance with the mortgage terms referred to in item 3 and the borrower(s) and every other signatory agree(s) to be bound by, and acknowledge(s) receipt of a true copy of, those terms.

Officer Signature(s)

Execution Date		
Y	M	D
04	01	07

 Borrower(s) Signature(s)

PETER W. DEMEO
Barrister & Solicitor
102 - 3930 Shelbourne St.
Victoria, BC, V8P 5P6

JOHN HENRY WILKINSON

PATRICIA MARY WILKINSON

BORROWER(S) (MORTGAGOR(S)): (including postal address(es))

STC? YES

Pick up STC?

3. BORROWER(S) (MORTGAGOR(S)): (including postal address(es) and postal code(s)) **Use Schedule**

JOHN HENRY WILKINSON, Dental Mechanic

After the clients are gone, Peter digitally signs the form and stares at the camera. (Note: camera and grin not required)




More than one "kind" of certificate can be used with Juricert and e-filing

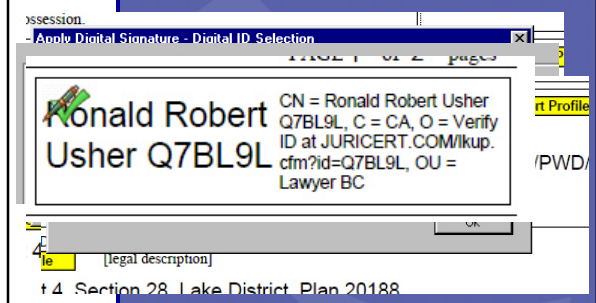


- Un-managed Acrobat certificate created by Juricert staff (only for use in Acrobat 6)
- Managed cert created by a Lotus Domino system (useable in email but root unknown)
- Fully Managed and recognized certificate from IdScript / DST-Identrus – root certificate known

The form is opened on his computer and the signing box is selected..



“signing” is just a few clicks



Roz then sends the mortgage to the LTB via BC Online..



I wonder if it will work...




Off it goes...to the network hub



Into the Fiber Optic Internet Connection...



**Off to 4300 Seymour in
Saanich, B.C.**



B.C. Government Data Centre

**Where Government
technology staff wait
anxiously..**



**Meanwhile, over at the Land
Title Branch
Headquarters...**



Staff try to stay calm..




Darcy and Denis check the time...




Did it work??

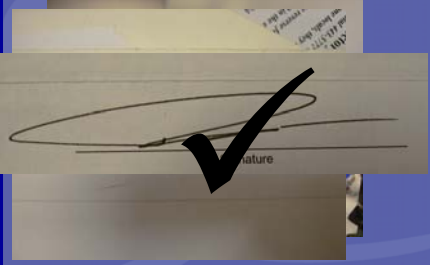





Meanwhile.....

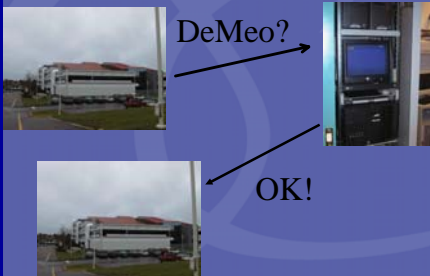


DeMeo's membership status and digital signature is checked by Juricert at the Law Society





(Actually the LTB's computer checks with the Juricert system)





Once the signature had been validated and other basic “edit” checks were completed by the LTB systems...the mortgage was accepted into the system...and electronically “marked up”




Huston, we don't have a problem! (at least with the mortgage...)



“numbers” and the stamped document are received in less than 2 minutes.

PAGE 1 of 2 pages

Rece
5)  Peter DeMeo IPP2VL 14

Digitally signed by Peter DeMeo IPP2VL
DN: CN = Peter DeMeo IPP2VL, C = CA, O
= Verify ID at www.jurioert.com/LKUP.cfm?
id=IPP2VL, OU = Practising Lawyer BC
Reason: Officer Certification for L.T.O.
purposes
Date: 2004.01.07 12:48:49 -0800

gent) **Import Profile** SUB BCOL ACCT: 995327

Vilkinson 33445/PWD/ab/rg

When printed the yellow boxes and the red "X" do not show

CA4



Document Copy

PAGE 1 of 2 pages

Peter DeMeo
IPP2VL

Digitally signed by Peter DeMeo (IPP2VL), DN: cn = Peter DeMeo (IPP2VL), o = CA, c = BC, email = peter.demeo@bc.ca, ou = BC, st = BC, serial = 13117613, version = 1.0, purpose = I am approving the document, Date: 2004.01.07 12:46:40 -0800

(or agent)

SUB BCOL ACCT:
995327

Wilkinson 33445/PWD/ab/rg

Fees taken from the firm's BC OnLine account are noted on the document

Fee Collected for Document: \$55.00

RTGAGOR(S): (including postal address(es) and postal code(s))

V8Z 3E9

STC Fee Collected for Document: \$9.00

Interest Rate:

(c) Interest Adjustment Y

The image was available from BC OnLine a few minutes later...



PENDING

VICA4

M39824 2004-01-07-13.06.54.416552

REGISTERED VICA4

M43508 2004-01-12-16.43.49.007359

LAND TITLE ACT
1 B (Section 225)

Received Land Title Branch: VIC
Jan-07-2004 12:49:49.001

CA4

So what did not happen in regard to this filing?



Off to the courier...



To the Agents reception area



Processed by the agents



Over to the LTO



Documents processed by the cashier



Agents wait for markup



Scanned and marked up



Back to the agents



Courier back to you




Handled by the receptionist again, and then



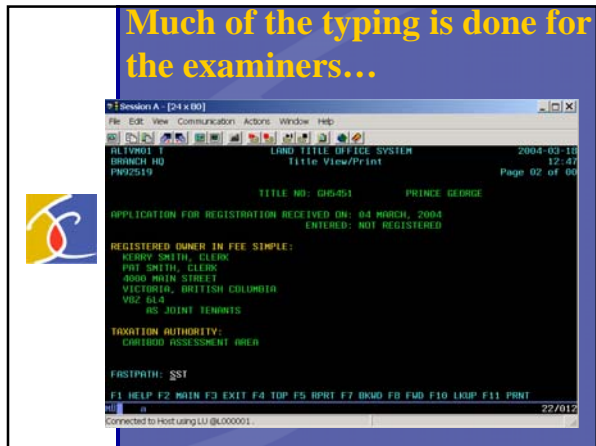
Back to the conveyancers



The major benefit to the LTO...



Much of the typing is done for the examiners...

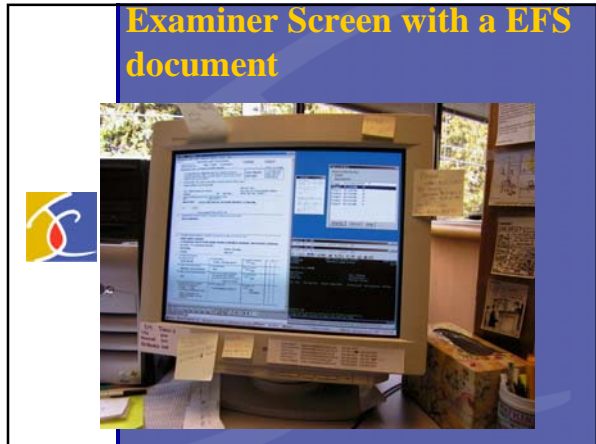


The image the examiner looks at is very clear...




LENDER(S) (MORTGAGEE(S)): (including
HSBC BANK CANADA
A CANADIAN CHARTERED BANK
AT #100 - 771 VERNON AVENUE
VICTORIA
Canada

Examiner Screen with a EFS document



On receiving the PDF File...



- Checks the form version
 - Checks for changes to the document post signing
 - Validates the signers status by comparing SHA1 Hash values
 - Ensures there is a valid property identifier ("PID")
- THEN
- Date, Time, and Number stamps the PDF
 - Extracts the data for posting to the DB2 Database
 - Creates a TIFF image for ImagePlus data base
 - Creates new "draft title" for the examiner
 - Processes fee payments
 - Reports to submitter

Success...



- From creation of the document to "stamped" and "bankable" document back in less than 15 minutes
- Courier and agents fees saved
- STC ordered with no extra paperwork
- Could immediately report to the bank
- Accounting details reported to the firm's bookkeeper (via email delivery of filing report)
- Both secretary and lawyer quickly able to finish the work on the file
- Approximately \$400 Million of Property Transfers and Mortgages filed so far
- Survived "DOS attack" by another system that went live at the same time for Corporation filings

We tried to find the "right"...



- Level of Security
- Technical Simplicity
- Balance of benefits (between gov. and others)
- Division of human labour and computer processing
- Allocation of costs (LTO/Juricert/Users)
- Fee Level (\$2.50 per digital signature)
- Extensibility (eg. digital survey plans)
- Amount of process change
- Respect for a system that was already "world class" for the government and consumers (and their lawyers)
- Platform for future refinement and development



Thank you for your time and interest. Please call if you have any comments, questions or concerns.

Ron Usher
The Law Society of B.C. / Juricert Services Inc.
845 Cambie St.
Vancouver, B.C. Canada V6B 4Z9
(604) 605-5310 rusher@juricert.ca
www.juricert.ca

PKI: Ten Years Later

Carlisle Adams
University of Ottawa
Ottawa, Canada
cadams@site.uottawa.ca

Mike Just
Treasury Board of Canada, Secretariat
Ottawa, Canada
just.mike@tbs-sct.gc.ca

Abstract

In this paper, we examine the history and evolution of so-called Public Key Infrastructure (PKI). We compare the original definition of PKI with a broader and more flexible definition that better reflects the variety of implementation philosophies available today. This current definition shows how the understanding of this technology has matured (although its essential characteristics have remained unchanged) and is derived, at least in part, from an evaluation and comparison of several quite different forms of PKI as well as a consideration of PKI criticisms over the years. The original definition of PKI may be dead or dying, but PKI technology continues to thrive as an extremely useful (and, in some cases, necessary) authentication solution.

1 Introduction

The technology known as “PKI” has been simultaneously maligned and praised. PKI praise can come in two flavours. The first results from a dislike for other security technologies. For example, a dislike for password-based authentication may result in a stronger preference for PKI solutions. Secondly, public-key technology offers some important benefits that are not similarly offered by other technologies, such as digital signatures. However, PKI is equally, if not more often, criticized. Difficulties around issues such as application integration, interoperability, and trust often lead critics to predict the end of PKI. While these shortcomings are very real, other issues have often been raised that either are orthogonal to PKI, or similarly impact non-public-key-based technology. In this paper, we try to identify such issues so that their true impact on PKI can be understood.

We attempt to provide some clarity to the status of PKI by discussing how it is understood today, compared with how it was initially defined. We examine the ways in which this updated definition encompasses ten years of PKI evolution, while remaining true to the original, essential characteristics of this technology. In order to do this, we compare several different PKI implementation models and see what lessons can be learned from some previous criticisms of PKI.

In Section 2, review and highlight several concepts related to public key technology and introduce six components that will contribute to our definition of a PKI. Section 3 reviews four PKI examples relative to these PKI components. In Section 4, we review and critique some well-known criticisms of PKI. Section 5 builds upon the previous PKI components, examples and critique to provide a modern definition for a PKI, while Section 6 recognizes those areas that still require development in order to support more successful PKI deployment.

2 Public Key Technology

In this section, we briefly examine some of the cryptographic properties of a PKI, and proceed to discuss how these properties may be used in practice.

2.1 Public Key Cryptography

Public key (a.k.a. “two key” or “asymmetric”) cryptography was invented by Diffie and Hellman in 1976 [DH76]. Unlike secret key (a.k.a. “symmetric”) cryptography, in which the same key K is shared between parties A and B, pairs of corresponding private and public keys for each user allow the unique realization of some operations. Specifically, let the respective private and public keys, $priv_A$ and pub_A , belong to party A. By operating on data with $priv_A$, A can *digitally sign* data that is verifiable by party B (or any other party) operating on the signed data with pub_A . Equivalently, party B (or any other party) can encrypt data using pub_A , where the encrypted data can only be decrypted with $priv_A$.

The true power of public key cryptography lies in the possession of a private key, *uniquely*, by each party. The “demonstration of knowledge” of the private key by operating on data with said key, provides a powerful tool that distinguishes asymmetric cryptography from its secret key counterpart.

2.2 Public Key Cryptography in Practice

Most secure transfers of data involve an exchange between identifiable parties. On its own, public key cryptography only supports asymmetric, mathematical operations on data; it does not by itself provide a connection to applications or environments such as e-commerce, e-mail, or the Web.

To provide such a connection, several additional pieces are necessary. These additional pieces form the definition of a PKI – an “infrastructure” that makes public key technology available to the applications and environments that wish to use it. In subsection 2.3 below, we identify several components that are integral to an infrastructure for supporting public key cryptography (these components are used to capture the evolving definition of a PKI in Section 5).

Identification is a property that is particularly critical to a PKI, and (at least historically) a strong differentiator between some different PKIs. Specifically, public key cryptography is made considerably more useful if the public key is *bound* to a so-called *identifier*. As distinguished below, this *identifier* may or may not provide direct information regarding an actual *identity*. This identifier may be an

- *Anonym*: “No name”; a single-use identifier providing no information as to the identity of the key owner.
- *Pseudonym*: “False name”; providing a “pretend” identity that can be used over a period of time to protect the real identity of the key owner.
- *Veronym*: “True name”; providing the identity of the key owner.

The identifier is typically meaningful only within a specific context or environment; it may or may not need to be globally unique, depending upon the applications for which it will be used.

Parties that use public key cryptography for encrypting data, or for verifying digitally signed data, will rely on the binding of an identifier to the public key (whether this binding is preserved in a certificate or database, for example) in order to associate that key with an entity with which they may have past or future interactions. This also supports repeated use of the same public key, whether or not the key is directly associated with an actual identity.

2.3 Public Key Infrastructure

Approximately ten years ago, the 1993 version of the ISO/IEC CCITT/ITU-T International Standard X.509 began to be disseminated, recognized, and implemented in small-scale environments. Late 1993 and early 1994 was effectively the beginning of PKI (although that

acronym had yet to be coined) because that version of the X.509 standard – more than the 1988 version – fleshed out some of the important details of certificates, certification authorities, and related concepts.¹

In those early days, a “PKI” was defined fairly rigidly, although with hindsight we can identify six major components to the definition that are still critical today, three to do with the *validity* of bindings (authority² functions), and three to do with the *use* of bindings (client functions). With respect to the validity of the binding between a public key and an identifier, what is needed is

1. an *authority* whose responsibility it is to create and destroy these bindings, as required, or aid in related authoritative actions,
2. an *issuance process* for expressing these bindings in a way that can be understood by other parties (i.e., in an agreed syntax) and for making this information available to parties that wish to know it, and
3. a *termination process* for breaking bindings when necessary and making this information available to parties that need to know it.

With respect to the use of such bindings, what is needed is

4. an *anchor management process* for augmenting or diminishing the set of authority public keys that will serve as roots or trust anchors for the client³,
5. a *private key management process* for ensuring that a client private key can be used for its desired purpose (this can include key pair generation and update, registering and binding an identifier to the corresponding public key, proper protection of the private key while it is valid, and backup & recovery of the private key in case of loss), and
6. a *binding validation process* for determining when the client should trust that a given public key (retrieved or acquired from some external entity) is authentically associated with a given identifier.

In Section 5, we develop a more detailed definition of a PKI, based on these components, that reflects the decade-long evolution of a PKI. The degree to which these components are implemented is commonly a risk management decision. The PKI examples in the next section can differ based upon such choices.

In the original Diffie and Hellman model [DH76], public keys would be retrieved from a secured repository. The security of this repository served to bind the public key to other attributes of the key owner. In support of offline binding production and distribution, Kohnfelder introduced the notion of a certificate [Kohn78], whereby a public key and an identifier (e.g., a name) were placed in a data structure signed by a Certification Authority (CA) and made available in an unsecured repository.

Various PKI systems can be distinguished and compared based upon the above six PKI characteristics. In Section 3, we categorize several examples of PKI systems with respect to these characteristics. It is particularly interesting to note that one of these examples makes use of

¹ Though the first version of Pretty Good Privacy (PGP) appeared in 1991, it wasn't until later that features consistent with a PKI were provided (see Section 3.2).

² An “authority” may be any specially designated entity, though an end-entity client may also act authoritatively.

³ These roots form the axiomatic elements of direct trust for a client. Trust in other public keys is derived from these roots.

Diffie and Hellman's original concept of a secured repository (AADS; see Section 3.3), while the remaining examples use "certificates" with varying syntax.

3 PKI Examples

Over the past 10-15 years, there have been several examples of public-key technology solutions. Below, we focus on those solutions that offer the best contrasts within the PKI components identified above. The representative example names (X.509, PGP, AADS/X9.59, SPKI) are quite overloaded with varying descriptions, as they may refer to several standards or even several varying implementations of those standards. In our review below, we have tried to focus on those features that are independent of specific product implementations yet representative of distinctive features for each PKI example.

On a related note, it is also recognized that over such a time period, the solutions have each grown and matured greatly. Though we attempt to identify this growth, our main purpose is to identify the philosophical differences between the solutions, so that not all features of each PKI solution may be acknowledged.

3.1 X.509

The X.509 standard [X509-00] and its Internet profile [RFC3280] do well to represent the PKI components identified in the previous section. In most cases, implementations differ based upon the rigour with which they implement the suite of appropriate standards (e.g., see the exhaustive list of Internet standards for X.509 [PKIX-WG]). Below, we examine relevant components of an X.509 PKI.

- *Authority.* A Certification Authority (CA) issues X.509 certificates that bind the public key to a Distinguished Name (DN) identifier (although other name forms are also allowed), in addition to other information contained in the certificate. An Attribute Authority (AA) is similarly defined, and binds more general attributes to one another in an attribute certificate, and provides an optional link to a corresponding public key certificate.
- *Issuance process.* Typically, though not necessarily, certificate issuance involves a Registration Authority (RA) responsible for registering the user (including their identification, if performed). Traditionally, the DN and alternative name forms would be veronymous. However, neither the ASN.1 syntax nor the standard restricts this so that anonymous or pseudonymous name forms are fully supported.⁴ Once issued by a CA, certificates require no further integrity protection and may be distributed amongst parties or made available in a repository. This repository is commonly an X.500 or LDAP directory, though various other repositories are typically supported now, including Web servers. Retrieval is predicated upon knowing the identifier of the certificate holder (typically the DN, although an email address contained in the certificate can also be used).
- *Termination process.* Certificates contain an expiry date that acts as a default termination date for the certificate. Certificates may also be "revoked" prior to their expiry, in which case the revocation information must be disseminated. There are traditionally two ways for this to be achieved: (i) by posting information regarding the revocation in a Certificate Revocation List (CRL), or (ii) by making the revocation information available through an Online Certificate Status Protocol (OCSP) responder [RFC2560]. The location of revocation information is typically included within the certificate that is being verified (e.g., as a URL for the CRL).

⁴ See [Just03] for an example of an X.509 PKI with a pseudonymous certificate identifier.

- *Anchor management.* The standards describe protocols for retrieving trust anchors as part of the registration (and key update) process(es). Depending upon the implementation, a client may be able to trust a number of trust anchors simultaneously (as part of a certificate trust list). Traditionally, there are two forms of trust for X.509 certificates. In the first, the application software holds the public key of a root CA. All certificates that may be trusted by this client are issued, either directly or indirectly (e.g., within a hierarchy), by this CA. In the second form of trust, a party holds the public key of the CA that issued their own certificate.
- *Private key management.* The standards support protocols for renewing key material prior to the expiry of the corresponding certificate. They also support the backup of decryption keys (and, more importantly, their recovery). The standards also allow a separate lifetime for the private key itself, primarily in support of preventing the creation of signatures too close to the time of certificate expiry, though this lifetime value is also helpful to trigger a timely key update in support of uninterrupted client operation.
- *Binding validation.* Clients use their trust anchors and possibly chain building to establish certificate trust. Trust in certificates issued by other CAs may be obtained through cross-certification between the CAs, or possibly by the party importing or retrieving the certificates of the other CAs as necessary. There are numerous variations to these two simple trust models. Traditionally, clients would be required to retrieve and validate the entire chain of certificates, though recent standards have been developed to support the off-loading of some of these operations [RFC3379].

In the often-cited “browser-based PKI”, it is important to recognize that certificates are issued to servers, while clients use those certificates to authenticate a server and establish a confidential communication channel.⁵ Clients retrieve the server’s certificate a part of the SSL protocol [Resc01]. The termination process supports expiry, though automated revocation support is minimal and inconsistent. Client anchor management is essentially static, and established by the version of Web browser being used, though users can manually update their trust anchors, if they so desire.

3.2 PGP

Though the first version of Pretty Good Privacy (PGP) appeared in 1991, its primary focus was in the support of public key cryptographic operations, not the provision of a PKI. Later versions supported notions such as key servers (see, for example, [PGPk]) thereby supporting an “infrastructure” for the management of public key information. In more recent times, PGP has highlighted its ability to also support features similar to X.509 [PGP99]. Traditionally, however, PGP has been distinguished by its distributed approach to key management. PGP certificates are formed of one or more “certifications”, which bind keys to user information with a digital signature. There is great flexibility in the key and user information that can be conveyed in a single certificate.

- *Authority.* Traditionally, a PGP PKI avoided the need for any form of authority. As discussed below, trust relies upon a “web” of users. However, the syntax does not preclude “authoritative users” that might act in a similar fashion to a Certification Authority (CA). For many communities, the lack of an authority greatly eases initial deployment as small communities need only rely upon the bilateral sharing of certificates among users who wish to communicate securely.
- *Issuance process.* Certificates are created and populated by the key owner. They can be distributed via email, Web pages, key servers, and/or other means. The identifier is typically an email address, though there is nothing to preclude other identifiers.

⁵ Though client authentication with certificates is supported by the standards, it is not often implemented.

- *Termination process.* Certificates can contain an expiry date (though no expiry date need be set) that acts as a default termination date for the certificate. The distribution of certifications is not managed, so that manual revocation would have to be performed, though revocation information can be made available in a similar fashion to publication of the certificate.
- *Anchor management.* Trust must be anchored in the *direct trust* of other users' certificates. Various mechanisms can be used to establish this direct trust base. For example, two users can exchange key information by email, but verify the authenticity of the exchange by exchanging message digests for the key information by phone. Such key information serves the role of PGP "roots" or "trust anchors."
- *Private key management.* Lacking a 3rd-party authority, private key management is the responsibility of the key owner. For example, key owners can backup private keys on a separate disk. It would be a simple task for software to remind users of the need to update their key material. Similar ease would allow an update of key material, since no communication with an authority is required, though updated key distribution would still be performed by the key owner.
- *Binding validation.* Based upon some initial, direct trust, there are a couple of options for indirectly extending trust to others. With hierarchical trust ("chain of trust"), you trust others that are trusted by people you trust. With cumulative trust ("web of trust"), you trust others only when they are trusted by a number of people you trust.

3.3 AADS/ ANSI X9.59

ANSI X9.59⁶ is a financial industry standard for secure financial payments, based on AADS (Account Authority Digital Signature) [AADS99]. For our purposes, we use it as an example of a non-certificate-based public key solution.

A public key is stored and managed as part of a key owner's financial account, along with other identity or authorization information. So, the issuer is an authority that already has a relationship with the user, and thus should be able to easily identify said user.

- *Authority.* The authority, or maintainer of the binding, is the user account manager. The public key serves, quite simply, as an additional attribute to a user's financial account record.
- *Issuance process.* Users may be identified by their account manager, based upon shared financial secrets, or other account information. The public key is retained by the manager. For AADS, the public key need only be accessed by an authentic request and response with the account manager to retrieve the public key for signature verification. For other applications, this could be easily adapted to allow for public-key encryption. Note that by not relying on a certificate, the AADS solution is more similar to the ideas of Diffie and Hellman than those of Khonfelder (see Section 2.3), except that with AADS the repository of public keys is built, held, and used by the relying party alone, whereas with the original Diffie-Hellman proposal this repository was to be created for the use of the whole world.
- *Termination process.* There are no "certificates" and use of any public key is always initiated by a request to the account manager. Expiry or revocation occurs by removing or replacing the public key for a user's account. Therefore, a type of immediate, online validation is supported as part of the public key retrieval.
- *Anchor management.* Relying parties require a method by which they can trust the account manager. A method similar to server-authenticated SSL would suffice. An initial trust anchor could be retrieved when the user's key pair is generated, for example.
- *Private key management.* Updates to private key material may be managed and initiated by the account manager. In the case of AADS, since only digital signature operations are

⁶ See <http://www.ansi.org/>

performed, there is no need to backup private key material. If encryption operations were supported with the user public keys, there may be a need for key backup support.

- *Binding validation.* Trust is isolated to the domain of a single account manager. As mentioned above, online trust validation is implicit with the retrieval of the public key.

Similar to SSL representing a client-server PKI implementation using X.509 certificates, SSH [SSH03] is representative of a certificate-less client-server PKI implementation. As part of the SSH transport protocol, clients establish trust in the server based on their trust in the server host key. Though there are options for how this trust might be established, including the client maintaining a store of trusted server host keys, or certification by a CA, SSH presents an interesting compromise whereby “the server name - host key association is not checked when connecting to the host for the first time” [SSH03]. Though introducing the potential for a middle-person attack, this novel variation offers great improvement for SSH bootstrapping. Once a server-authenticated, confidential channel is established, the client may authenticate to the server; this is often performed using password authentication.

3.4 SPKI

Simple Public Key Infrastructure [RFC2692, RFC2693] was developed in response to several criticisms of X.509. The major philosophical objection to X.509 surrounds its relation to X.500 naming. SPKI, more correctly an authorization infrastructure, relies upon the uniqueness of the combination of a pseudonym and a public key.

- *Authority.* SPKI focuses on the issuance of authorization information within certificates. Thus, an SPKI authority might be referred to as an authorization authority. With regard to an issuance authority, SPKI theory indicates that certificates may be generated “by any keyholder empowered to grant or delegate the authorization in questions.” [RFC2692]
- *Issuance process.* In support of the authorization information, the SPKI certificate syntax uses an S-Expression, which is a LISP-like expression using parentheses. *Authorization certificates* bind authorization information to a key, while *attribute certificates* bind an authorization to a name. The use of names differs from the initial use of global names for an X.500 directory, as part of X.509, and was inspired by the use of SDSI’s [SDSI96] local names. Combined with the (globally unique) hash of a public key, such a name can become globally unique.
- *Termination process.* Certificate lifetime is parameterized by a validity period so that certificates can be set to expire. Several options for certificate revocation are supported, including Certificate Revocation Lists (though they are not the preferred choice). Options for online revocation status checking are also supported. Preference is given to “positive statements” on certificate validity, so that a protocol returning an indication that a certificate is currently valid is favourable to one that returns a notice of invalidity.
- *Anchor management.* Details regarding anchor management are left open for developers so that, for example, protocols similar to those previously described could be used. For validation of authorization information, however, the relying party maintains an access control list (ACL).
- *Private key management.* The management of private keys depends upon the certificate issuer regarding issues of key backup; however, when used only for authorization purposes, the need for key backup is limited. Support for key updates does not appear to be standardized, so would be dependent upon the specifics of a particular implementation.
- *Binding validation.* The main difference with traditional X.509 is the use of the pseudonym for SPKI. Processing decisions for SPKI certificates are defined through “tuple reduction.” [RFC2693].

3.5 Summary of Examples

The following table compares the solutions based upon the *validity of bindings* (see Section 2.3).

PKI Solution	Authority	Issuance Process	Termination Process
X.509	Certification Authority (CA) Attribute Authority (AA). The CA is the owner / definer of the namespace for the identifier.	ASN.1 syntax Traditionally available from X.500 or LDAP directories.	Certificate contains an expiry date. Revocations posted through revocation lists, or made available through an OCSP responder.
PGP	No external authority required. Key pair and certificate are self-generated. The user (end entity) is the owner / definer of the namespace for his/her identifier.	Made available to others by key owner (e.g. via Web page, email signature, or key server).	Certificates can expire. Termination performed by key owner. Dissemination of termination notice by key owner as with certificate publication.
AADS/ X9.59	User account manager. The relying party (the account manager) is the owner / definer of the namespace for the identifier (the acc't. #).	Public keys available in secured repository from account manager.	Public keys removed from repository when binding is terminated.
SPKI	No explicit authority is required as the authorization granter or delegator may issue certificates. The relying party is the owner / definer of the namespace for the identifier.	Issue authorizations based on pseudonymous identifier or SDSI names.	Similar to X.509, though "positive statements" through online validation are preferred.

The following table compares the solutions based upon the *use of bindings* (see Section 2.3).

PKI Solution	Anchor Management Process	Private Key Management Process	Binding Validation Process
X.509	Single or multiple roots. Standardized protocols support changes.	Standardized protocols support update, backup and recovery.	Client search of Directory for cross certificates. Delegated path discovery and validation services are being standardized.
PGP	Direct trust of other user certificates. Trust anchor is user's own key(s).	Manual update, backup, and recovery performed by user.	Chain of trust, or web of trust.
AADS/ X9.59	Trust in account manager is required.	Depends upon expiry policy. Backup and recovery not a concern when only digital signatures are used.	Only direct validation through trusted key retrieval.
SPKI	Open to developer. Trust anchor is the ACL at the relying party.	Open to developer. Fewer backup and recovery requirements when certificates used only for authentication or authorization.	Tuple reduction.

These fundamental PKI examples contribute to a greater understanding of the different options available for PKIs within what is mistakenly viewed as a “rigid” structure. In the following section, we further examine criticism of PKI, identifying those issues that are specific to the components of this infrastructure. In Section 5, we use the examples and the lessons learned from the criticism to capture the evolutionary definition of PKI.

4 Criticism of PKI

Over the past ten years, PKI has been the subject of criticism from various quarters. Some of this criticism has been beneficial, driving the evolution of this technology and leading to a deeper understanding and broader application of PKI. However, much of the criticism has been misdirected, aimed at PKI when the actual problem or challenge is either independent of this technology, or common to many technologies.

In this section we review some popular PKI criticisms to see which can fairly be applied to the current state of the art. While it is certainly not the only collection of criticisms, arguably the best known collection can be found in the paper by Ellison and Schneier [EISc00]. We therefore use that paper as the basis for our examination of PKI, circa 2004.

“Ten Risks of PKI: What You’re not Being Told about Public Key Infrastructure” aims to explore some basic questions around PKI (“What good are certificates anyway? Are they secure? For what?”) so that potential users of this technology can be made aware of some of the risks involved with its use. This is unquestionably a worthy goal and will serve the industry well, but only if the highlighted risks are accurate and fair (i.e., legitimate criticisms of PKI technology). Let us examine some of the risks discussed in that paper.

Risk #1 (“Who do we trust, and for what?”) warns that the certificates issued by a CA should not be automatically trusted for a plethora of application-level purposes, such as making a micropayment or signing a million-dollar purchase order (“Who gave the CA the authority to grant such authorizations? Who made it trusted?”). Unfortunately, this criticism highlights only the misuse of PKI by some implementers; it is not a valid criticism of PKI itself. PKI is an authentication technology; authorization is an independent matter and may or may not be linked to authentication in any way. The authors suggest that “Many CAs sidestep the question of having no authority to delegate authorizations by issuing ID certificates.” However, the issuance of ID certificates is the primary function of a CA (not a “sidestep”). In some environments, it may be natural for information other than an identifier to be linked to the public key by the CA; for such situations a variety of authorities for such information may be used in conjunction with the CA (these are discussed as part of the evolving PKI definition in Section 5). On a related note, while certificate policies appear to contain some notion of authorization, they are more properly viewed as statements regarding the “quality” of the key. For example, given the specific process used to generate the key pair, the rigour with which identification of the key holder was done, the care with which the private key will be safeguarded, and so on, the CA declares (by including a policy to this effect in the certificate) that the public key can be used for signing million-dollar purchase orders. But this is not a granting of authority. In a properly-implemented system, the signer must still prove that s/he is authorized to sign such a purchase order (and this authorization will typically come from some entity in the environment that is not the CA). Certificate policy may be viewed as a “fit for purpose” declaration: if the signer is allowed to sign such a transaction, then this key pair can be used to create and to verify that signature.

Risk #2 (“Who is using my key?”) warns that the private key stored on your computer may not be secure (without physical access controls, TEMPEST shielding, air-gap network security, video surveillance, and so on). Clearly this is true, but is equally true of all technologies that store data on a computer. In order to address this, PKI has evolved to support both “soft token” solutions (in which the user retains the private key) and roaming solutions (in which the private key may be stored at a server). As always, there are security / convenience trade-offs for each. When stored at the user’s computer, the user can authenticate with his/her private key to a server that has an authentic copy of the public key. This is arguably more secure than solutions that store either a clear-text or hashed version of a password at a server in support of password-based authentication (the latter is susceptible to brute-force attack). The discussion about PKI vendors “lobbying for laws to the effect that if someone uses your private signing key, then you are not allowed to repudiate the signature” does not reflect any of the myriad debates we have heard and read on this topic. (On the contrary, if there is any reasonable evidence that someone else has used your private key, this is precisely when you can repudiate a digital signature.) In any case, in recognition of the vulnerabilities associated with typical computing platforms, PKI has come to strongly support alternative devices, such as hardware tokens and smart cards, for storing private keys and trust anchors.

Risk #3 (“How secure is the verifying computer?”) examines the insecure computer question (i.e., Risk #2) again, but this time from the side of the verifying machine. As above, this risk is shared by all technologies that use computers and is not specific to PKI. Again, alternative storage devices can be helpful here.

Risk #4 (“Which John Robinson is he?”) warns that the name in a certificate may not be as valuable as it appears to be (“You may know only one John Robinson personally, but how many does the CA know? ... How many John Robinsons are in the New York City phone book, much less in a hypothetical phone book for the global Internet?”) Additionally, the authors ask, “How do you find out if the particular John Robinson certificate you received is your friend’s

certificate?” But one could equally ask how you find out if the e-mail you received is from your friend John Robinson or from some other John Robinson, or how you find out if the person on the other end of the telephone is your friend John Robinson, or how you find out if the postcard you received in your mailbox is from your friend John Robinson. Real life requires us to be able to resolve the potential ambiguity with regard to a name, and we do this all the time, but this is not a problem that is either created, or purported to be solved, by PKI. Users in the electronic world, as in the physical world, need to be able to do the mapping between name and identity whether or not PKI was ever invented. This is true of all authentication technologies. A PKI binds an identifier to a public key. Associating that identifier with an identity, or with entitlements within the context of the application being used, is outside of the scope of PKI; it always has been. Applications that rely upon PKI for authentication need to recognize that this issue is not solved by PKI. (More discussion of this topic can be found in Chapter 14 of [AL03]).

Risk #5 (“Is the CA an authority?”) warns that the CA may not be an authority on what the certificate contains (e.g., the corporate name of the keyholder and the DNS name of the server in an SSL server certificate). There are authorities on DNS name assignments, and authorities on corporate names, but the CA is likely to be neither of these. This is quite true but, as stated above, it is not necessarily the job of the CA to create names, or even assign names to entities; its primary function (and its authority) is to bind an identifier to a public key. As time has passed, it has become more generally recognized that a CA may make use of other authorities in order to do this task. In particular, it will often collaborate with other naming authorities to ensure that the information in a certificate is as accurate as possible.

Risk #6 (“Is the user part of the security design?”) warns that users will often make security decisions (such as whether to shop at a given SSL-protected Web page) without even seeing the certificate involved or knowing whether it has any relation to what is displayed. This is certainly true, but is equally true of many security technologies. If a security infrastructure provides a technical means by which application security decisions can be automated and enforced, and then these means are not used, this is not the fault of the infrastructure and is not a risk of the infrastructure. More accurately, the “risk” is that security software is often not implemented correctly or used properly, but this is true of all security software, everywhere, and has nothing specific to do with PKI. However, in general, PKI implementations do need to provide for more useful interaction with the user [WhTy99].

Risk #7 (“Was it one CA or a CA plus a Registration Authority?”) warns that “the RA+CA model allows some entity (the CA) that is not an authority on the contents to forge a certificate with that contents.” This is true, but authorities in any system can always abuse their power. This is not a risk specific to PKI, but is true for all systems, everywhere. Furthermore, even if an RA+CA combination can be less secure than a single CA, there are certainly environments in which placing all power into the hands of a single authority can also be a highly risky thing to do. As in any system, it is important to choose the authorities carefully or the system will not work as intended. Over the years, explicit statements of CA practices and policies (see, for example, the framework specified in [RFC3647]) have come to be used throughout the PKI community so that external auditors and inspectors can check whether a given CA is abusing its power in some way.

Risk #8 (“How did the CA identify the certificate holder?”) warns that the CA may not use good information to check the identity of the entity applying for the certificate, or may not ensure that this entity really controls the private key corresponding to the public key being certified. This is true, but again is not a risk specific to PKI. If authorities in any system do not do their jobs diligently and with integrity, the system cannot be expected to work. This is not a failing of the system itself. Authorities in a PKI (in particular, the CAs) need to be chosen carefully and

trained well, but this is no different from any other system. As in Risk #7 above, auditable statements of CA practices and policies can be helpful to avoid problems in this area.

Risk #9 (“How secure are the certificate practices?”) warns, in a nutshell, that “Certificates must be used properly if you want security.” It is hard to imagine that anyone would argue with such a statement. But passwords must be used properly if you want security; biometrics must be used properly if you want security; smart cards must be used properly if you want security; and so on. Once again, this is not a risk specific to PKI; this basic statement holds true for all security technologies.

Risk #10 (“Why are we using the CA process, anyway?”) warns that PKI does not solve all security problems, even though it is sometimes marketed and sold under that premise. This is in some ways a fair criticism, as some over-zealous marketing executives have sought to increase profits by stretching the truth in several directions. However, this is not a risk of PKI. All this highlights is a need to get accurate information out regarding what PKI actually is, and what it actually does. Things have improved significantly in this area in the past few years, but more can certainly be done.

In summary, we find that of the popular PKI criticisms voiced in the literature and at various conferences, many do not apply to PKI at all, and most of the rest apply equally to all security technologies. (As a measure of items related to implementing and deploying a PKI, however, they do highlight some specific concerns. And, as with the other security technologies to which they pertain, solutions – often outside the scope of a PKI – can be applied.) For the remaining criticisms that are accurate and valid, the evolution of this technology has come to understand and address these comments so that the current (at least theoretical) view of PKI no longer appears to be deficient in these ways. Such criticisms have therefore been very beneficial to the industry as a whole.

5 PKI Evolution and a Current Definition

The comparison of approaches in Section 3 makes it clear that a PKI can be instantiated today in many different ways. But ten years ago, several of the above instantiations would not have fit into the “accepted vision” of what a PKI was. Clearly, something has changed, but is it the essence of the definition, or the implementation details? Guided by the general definition of Section 2.3, we see that the essence remains intact; only our understanding of each of the components of that definition has evolved over time.

Definition 1994. In 1994, the six components of the general definition given in Section 2.3 were restricted in the following ways.

1. *Authority.* The authority was always and only a Certification Authority (CA). There was no place in the PKI for any other kind of authority.
2. *Issuance process.* The syntax was always and only an X.509 public key certificate which binds a public key to a Distinguished Name (DN) for the user. The certificate was made available to other parties through the use of an X.500 Directory.
3. *Termination process.* The termination process was always and only a Certificate Revocation List (CRL) which could be made available to other parties through the use of an X.500 Directory (perhaps pointed at using a CRL Distribution Point in the certificate).
4. *Anchor management process.* A user may trust the CA that is “closest” to him/her (i.e., the CA that actually issued the user’s certificate) or may trust the root of a hierarchy of CAs

which includes the CA that actually issued the user's certificate. In either case, however, the client code was pre-installed with a single trust anchor and no changes were possible.

5. *Private key management process.* Very little of this was specified, although it was generally assumed that key generation occurred at the CA, registration occurred via an out-of-band, in-person process, and private keys were "safe" in the user's local environment (perhaps protected by a password).
6. *Binding validation process.* Client machines had to be configured with a large, special-purpose software toolkit that could understand all the details of certificate processing and could make validated public keys available to application code.

We now propose an updated definition of PKI.

Definition 2004. By 2004, after ten years of evolution that has resulted from extensive discussion, research, and implementation by various interested parties around the world, we find that each of the above six components of the definition has broadened considerably. However, interestingly, the same six components comprise the core of the definition. That is, the essential characteristics of the definition remain unchanged, even if the thinking about how to realize these characteristics has deepened and matured over time.

1. *Authority.* The notion of an "authority" has broadened from the CA that is "closest" to a user, to a CA that may be "farther away" (e.g., at the root of the user's hierarchy), to a CA that may be even farther away (e.g., at the root of a different hierarchy in another domain), to a CA that may be entirely independent of the user (e.g., one offered as a public service). In addition, the authoritative role of a CA might be performed by an end entity. Furthermore, it is now recognized that a CA may make use of other entities prior to issuing a binding. For example, an *Identification entity* (perhaps a Registration Authority, or some other entity altogether, such as a Naming Authority) may be used to properly determine the correctness of an identifier (on behalf of, or at the request of, a CA) before the identifier is bound to the public key. As well, PKI now recognizes the utility and value of other authorities in the environment that are not CAs, such as OCSP Responders, certificate path validation authorities, Attribute Authorities, and so on.
2. *Issuance process.* A number of different syntax proposals have been discussed and implemented over the years, and it is now well recognized that some environments will be more suited to a particular syntax than others. There is therefore a need for various ways of encoding the binding expressed by an authority. Similarly, options for the type of identifier (see Section 2.2), and the actual location of the trustworthy binding, have evolved as different choices. Certificate formats such as PGP, SPKI, SAML, XKMS Responses (see Section 6), and so on, all have a place in this broader definition. Furthermore, it is recognized that X.500 Directories are but one possible mechanism for making these bindings available to other entities, and many other technologies are now commonly in use to achieve this.
3. *Termination process.* Breaking the binding between a public key and an identifier can now use many more mechanisms than the traditional CRL. Online certificate status checkers (e.g., OCSP) were an early step in this direction, but even broader online services, such as delegated path validation [RFC3379] and XKMS servers, have also been envisioned and implemented. The use of on-line checking includes the option of online certificate retrieval, where only if the certificate is available, is it considered valid at the time of retrieval. The PKI community has come to realize that the information regarding a revoked binding needs to take different forms and use different delivery mechanisms for different environments.
4. *Anchor management process.* In support of the broader definition of authority, mechanisms for establishing how different parties accept the bindings issued by an authority have been defined and used, including trust root installation, cross-certification, trust lists, and so on. It

has become recognized that the trust anchors for a user will typically be a set of size greater than one, and that the members of this set will need to change over time.

5. *Private key management process.* Thinking in this area has broadened significantly over the past ten years as the need for flexibility in different environments became clear. To list a few examples, key generation may take place at the CA, at the client, or at a third party and protocols have been developed to handle any of the three options securely (along with protocols to allow secure backup and recovery of key material). Registration might need to be an online process (rather than an offline process) for efficiency and user convenience. As well, private keys might be stored in software, in hardware tokens, in smart cards, or in other formats, and might be stored with the user or at some 3rd-party server; protocols and interfaces have been developed over the years to handle all of these options.
6. *Binding validation process.* With respect to software, there was a growing concern that large, special-purpose toolkits were not the best alternative for some environments (for a number of reasons, including cost, size requirements, and complexity of upgrades). Interest in this area shifted to so-called “thin clients”: toolkits that were very small for fast and easy download (e.g., in Java applet form). But there was also a growing realization that, in some environments at least, the ideal situation would be native applications (e.g., off-the-shelf browsers) that could properly understand all the details of certificate processing.

The current view of PKI, as expressed in “Definition 2004”, is a reflection of the evolution that has occurred in this community over the past ten years. It benefits from the innovative thinking and fruitful technical discussion of researchers the world over, and has been steered greatly in more practical and useful directions by constructive criticism and numerous implementation efforts (see Sections 3 and 4 above). This definition, we believe, represents the “state of the art” in the understanding of PKI.

6 Moving From Theory to Practice

Reflecting – in an updated definition – the evolution (and the occasional revolution!) that has occurred over the years may be a useful step, but it is not sufficient. Clearly, this deeper understanding of PKI needs to be embraced in a real way in real implementations. This is not to suggest that a given PKI implementation should strive to be all things to all people. If we as a community have learned anything over the past decade, it is that the many options available for each component of the definition preclude any “one size fits all” PKI. However, even for a given set of choices, most (perhaps all) implementations can improve in both correct operation and suitability to a given environment. Many common implementation bugs and challenges have been summarized well by Gutmann [Gut, Gut02]. Specifically, Gutmann identifies issues regarding hierarchical naming, revocation, and certificate chain building. Current and prospective implementers of PKI technology would do well to look through some of this material.

One important realization of Gutmann is that original (and even some current) PKI implementations would “constrain the real world to match the PKI”, as opposed to “adapt[ing] the PKI design to the real world.” [Gut02]. It is hoped that we similarly capture this concern in our discussions above. In considering further issues that may remain regarding the deployment of PKI within some environments, a survey was recently performed by the OASIS PKI Technical Committee [OAS03]. The main impediments cited were the cost and lack of PKI support within client applications. Noteworthy in this regard are more recent PKI-related efforts that were motivated to address this specific concern. In particular, XML Key Management Services [XKMS03] have primarily been designed in order to abstract away some of the technical PKI detail in order to work with relatively simple clients. The generic protocol description for XKMS

should allow it to support any of the PKI examples discussed in Section 3. Key management issues are part of the key registration service component (X-KRSS), while key validation issues are part of the key information service component (X-KISS). A common Web services interface may go some way to aid the otherwise difficult process of integrating these components with existing software.

An area that is yet to be widely embraced in real implementations concerns the nature of the identifier used in a certificate. There are times when there is a legitimate need for this identifier to be veronymous, other times when a pseudonym would be preferable, and still other times when an anonym should be used (even within a single environment). Yet existing CAs are typically built to use only a single type of identifier (perhaps, if the CA is very flexible, in a range of formats). Standards, in their language and in their syntax, do not generally preclude the use of different identifier types, but history and tradition have made rigid interpretations and resulted in PKI deployments that are almost exclusively one type or another. More flexibility in this area (i.e., CAs that can bind keys to any of the three types, as required) would make PKIs more suited to many real-world requirements.

The goal of this paper has been to demonstrate that the PKI community has significantly broadened its understanding of this technology over the past ten years. The challenge now is to translate that understanding to real PKI implementations that solve authentication challenges in real, heterogeneous environments.

Acknowledgements

The authors wish to thank the reviewers for their numerous comments. They greatly improved the content and readability of this paper.

7 References

- [AADS99] L. Wheeler, "Account Authority Digital Signature and X9.59 Payment Standard", slides presented at the 3rd CACR Information Security Workshop, June 1999. (<http://www.cacr.math.uwaterloo.ca/conferences/1999/isw-june/wheeler.ppt>)
- [AL03] C. Adams, S. Lloyd, *Understanding PKI: Concepts, Standards, and Deployment Considerations, Second Edition*, Addison-Wesley, 2003.
- [DH76] W. Diffie, M. Hellman, "New Directions in Cryptography", *IEEE Transactions on Information Theory*, Vol. 22, No. 6, November 1976, pp. 644.
- [ElSc00] C. Ellison, B. Schneier, "Ten Risks of PKI: What You're not Being Told about Public Key Infrastructure", *Computer Security Journal*, vol.XVI, no.1, 2000.
- [Gut] P. Gutmann, "Everything you Never Wanted to Know about PKI but were Forced to Find Out"; see <http://www.cs.auckland.ac.nz/~pgut001/pubs/pkitutorial.pdf>
- [Gut02] P. Gutmann, "PKI: It's Not Dead, Just Resting", *IEEE Computer*, August 2002; see <http://www.cs.auckland.ac.nz/~pgut001/pubs/notdead.pdf>
- [Just03] M. Just, "An Overview of Public Key Certificate Support for Canada's Government On-Line (GOL) Initiative", to appear in *Proceedings of 2nd Annual PKI Research Workshop*, April 2003.
- [Kohn78] L. Kohnfelder, "Towards a Practical Public-key Cryptosystem", *MIT Thesis*, May, 1978.
- [OAS03] P. Doyle, S. Hanna, "Analysis of June 2003 Survey on Obstacles to PKI Deployment and Usage", report of the *OASIS PKI Technical Committee*, v1.0, 8 August 2003. Available at <http://www.oasis-open.org/committees/pki/pkiobstaclesjune2003surveyreport.pdf>.

- [PGP99] J. Callas, "The OpenPGP Standard", slides presented at the 3rd CACR Information Security Workshop, June 1999.
(<http://www.cacr.math.uwaterloo.ca/conferences/1999/isw-june/callas.ppt>)
- [PGPkS] The MIT PGP Key Server; see <http://pgp.mit.edu/>
- [PKIX-WG] IETF Public-Key Infrastructure X.509 (PKIX) Working Group; see <http://www.ietf.org/html.charters/pkix-charter.html>
- [Resc01] E. Rescorla, *SSL and TLS: Designing and Building Secure Systems*, Addison-Wesley, 2001.
- [RFC2560] M. Myers, R. Ankney, A. Malpani, S. Galperin, C. Adams, "X.509 Internet Public Key Infrastructure: Online Certificate Status Protocol – OCSP", Internet Request for Comments 2560, June 1999.
- [RFC2692] C. Ellison, "SPKI Requirements", *Internet Engineering Task Force (IETF) Request for Comments (RFC) 2692*, September 1999.
- [RFC2693] C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylönen, "SPKI Certificate Theory", *Internet Engineering Task Force (IETF) Request for Comments (RFC) 2693*, September 1999.
- [RFC3280] R. Housley, W. Polk, W. Ford, D. Solo, "Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List (CRL) Profile", Internet Request for Comments 3280, April 2002.
- [RFC3379] D. Pinkas, R. Housley, "Delegated Path Validation and Delegated Path Discovery Protocol Requirements", Internet Request for Comments 3379, September 2002.
- [RFC3647] S. Chokhani, W. Ford, R. Sabett, C. Merrill, S. Wu, "Internet X.509 Public Key Infrastructure: Certificate Policy and Certification Practices Framework", Internet Request for Comments 3647, November 2003.
- [SDSI96] R. Rivest, B. Lampson, "SDSI – A Simple Distributed Security Infrastructure", 17 September 1996, <http://theory.lcs.mit.edu/~rivest/sdsi10.html>
- [SSH03] T. Ylonen, D. Moffat, "SSH Protocol Architecture", *Internet Draft*, October 2003. Available at <http://www.ietf.org/internet-drafts/draft-ietf-secsh-architecture-15.txt>.
- [WhTy99] A. Whitten, J.D. Tygar, "Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0.", in *Proceedings of the 9th USENIX Security Symposium*, August 1999.
- [X509-00] ITU-T Recommendation X.509. *Information Technology – Open Systems Interconnection – The Directory: Public Key and Attribute Certificate Frameworks*. March 2000 (equivalent to ISO/IEC 9594-8:2001).
- [XKMS03] P. Hallam-Baker, "XML Key Management Specification", W3C Working Draft, v2.0, 18 April 2003. Available at <http://www.w3c.org/2001/XKMS/>.

PKI: Ten Years Later

Carlisle Adams

*School of Information Technology and Engineering
University of Ottawa*

Mike Just

Treasury Board of Canada, Secretariat

Outline

- Motivation
- Public key technology and PKI
- PKI examples
- PKI criticisms
- PKI evolution and a current definition
- The road ahead...

Motivation

- We have reached an anniversary in PKI
- Has our understanding of this technology grown in any way? If so, how?

PK Technology and PKI

- Public key technology
 - Each entity in a collection has a pair of keys
 - Alice has $pub_A, priv_A$
 - Enc, d-sig. possible (mathematical operations)
- Public Key Infrastructure (PKI)
 - Makes PK technology available to applications and environments that wish to use it
 - Enc, d-sig. possible (security operations)
 - Key pair bound to an entity identifier in a way that makes it useful to a variety of apps

PKI (cont'd)

- "Identifier"
 - Uniquely specifies entity within some context or environment (no ambiguity), but need not reveal actual identity
 - Anonym (single-use identifier; no mapping to entity)
 - Pseudonym (multiple-use identifier; no mapping to entity)
 - Veronym (multiple-use identifier; clear mapping to entity)
 - Context/environment need not be global in scope (depends on apps that will use keys)

PKI (cont'd)

- Binding of key pair and identifier
 - Validity of bindings
 - Authority (making & breaking)
 - Issuance process (syntax & dissemination)
 - Termination process (alerting)
 - Use of bindings
 - Anchor management process (augment & diminish)
 - Private key management process ("fit for purpose")
 - Binding validation process (trusting someone else's key)

Outline

- Motivation
- Public key technology and PKI
- **PKI examples**
- PKI criticisms
- PKI evolution and a current definition
- The road ahead...

PKI Examples

- Over the past ten years, there have been several different approaches to modeling and implementing a PKI
- These approaches can be compared based on the 6 components of the "binding" concept
- We look at the following:
 - X.509, PGP, X9.59, SPKI

Sample Comparisons

(see paper for others)

PKI Solution	Authority	Issuance Process
X.509		
PGP		
AADS / X9.59		
SPKI		

Sample Comparisons

(see paper for others)

PKI Solution	Authority	Issuance Process
X.509	CA, AA. CA is owner / definer of namespace.	ASN.1 syntax. X.500 or LDAP directories.
PGP	No external authority. User is owner / definer of namespace.	BNF syn. Issued by key owner (e.g., Web page, e-mail sig., key server).
AADS / X9.59	User account manager. Acct. mgr. is owner / definer of namespace.	(Raw) public keys available in secured repos. from acct. mgr.
SPKI	Authorization granter. Relying party is owner / definer of namespace.	S-expression syntax. Issued based on SDSI names or pseudon. Ids.

PKI Criticisms

- Many criticisms have been leveled at this technology
- Probably the best-known collection is the "10 Risks" paper by Ellison & Schneier
- But criticisms cannot always be taken at face value: need to consider whether the "flaw" being criticized is actually related to PKI or not

PKI Criticisms (cont'd)

- Examples:
 - Authentication versus authorization
 - Security of computing platforms
 - Linkage between identifier and real entity ("John Robinson problem")

PKI Criticisms (cont'd)

- **Understatement alert:** PKI has had its share of critics over the years
- A number of criticisms have been unjustified, and a number have been misdirected (aimed at PKI when the actual problem is elsewhere)
- The remainder have been very beneficial, driving evolution and leading to a deeper understanding of this technology

Outline

- Motivation
- Public key technology and PKI
- PKI examples
- PKI criticisms
- **PKI evolution and a current definition**
- The road ahead...

Evolution

- Ten years ago, the 1993 version of the ISO/IEC CCITT/ITU-T IS X.509 began to be disseminated, recognized, and implemented in small-scale environments
- Late 1993 / early 1994 was effectively the birth of PKI (although the acronym was yet to be coined)
 - Infrastructural considerations were paramount (how to make PK technology available to a wide variety of applications)

Evolution (cont'd)

- Initial definition (1994)
 - Authority: always and only a CA
 - Issuance: X.509 syntax; DN; X.500 Directory
 - Termination: CRL; X.500 Directory
 - Anchor: root of CA hierarchy
 - Private key: CA gen.; OOB reg.; local storage
 - Validation: large, special-purpose s/w toolkit

Evolution (cont'd)

- After ten years of extensive discussion, research, and implementation by numerous interested parties world-wide:
 - Each of the 6 components has broadened quite considerably with deeper understanding
 - BUT, the same 6 components comprise the core of the definition (i.e., the essential characteristics of the definition remain unchanged)

Evolution (cont'd)

- Current definition (2004)
 - Authority: multiple choices (incl. end entity)
 - Issuance: multiple choices (syntax & dissem.)
 - Termination: multiple choices (incl. online)
 - Anchor: multiple choices (augment & diminish)
 - Private key: multiple choices (gen., reg., storage)
 - Validation: mult. choices (thin client; native apps)

Outline

- Motivation
- Public key technology and PKI
- PKI examples
- PKI criticisms
- PKI evolution and a current definition
- The road ahead...

Future of PKI

- Moving from theory to practice
 - Over ten years, innovative thinking, fruitful technical discussion, constructive criticism, and implementation efforts have driven the recognition of the need for options
 - Research into secure architectures and secure protocols have made options possible
 - BUT options have yet to be embraced in a significant way in real products

Future of PKI (cont'd)

- Example: identifier bound to public key
 - Sometimes there are valid reasons for the identifier to be a veronym; sometimes a pseudonym; sometimes an anonym
 - Standards (in their language and syntax) do not preclude different identifier types
 - However, history and tradition have made rigid interpretations: PKI deployments are almost exclusively one type or another
 - WHY NOT HAVE CAs THAT CAN BIND KEYS TO ANY OF THE THREE TYPES, AS REQUIRED?
 - This would make PKIs more suited to real-world needs

Conclusion

- The goal of this work has been to demonstrate that the PKI community has significantly broadened its understanding of this technology over the past ten years
- The challenge now is to translate that understanding to real PKI deployments that solve authentication challenges in real, heterogeneous environments

An Examination of Asserted PKI Issues and Proposed Alternatives

John Linn, RSA Laboratories, Bedford, MA, USA
Marc Branchaud, RSA Security Inc., Vancouver, BC, Canada

15 March 2004

1 Introduction

Since the 1980s, public-key infrastructures (PKIs) have been widely anticipated as a primary means to make entities' keys available to others in a trusted fashion, thereby enabling a qualitative improvement in the protection and assurance of communications and transactions carried out over the Internet. Certificate-based authentication has become common practice in certain contexts, particularly in conjunction with SSL-protected web sites. In recent years, however, many commentators have lamented the fact that PKI has not achieved more pervasive adoption and deployment. Some, like [Clar01], [ElSc00], and [Gutt02], have concluded that PKI is a failure or does not address users' primary security needs. Opinions differ on the reasons for these results, but most can be distilled into a few general categories:

- A belief that demand for the services offered by PKI, in terms of PKI-integrated applications and/or security-oriented use cases for those applications, has not yet emerged to a degree sufficient to motivate deployment of a trust infrastructure.
- A belief that characteristics of current PKI architectures and implementations make them unnecessarily difficult to deploy, and/or that those characteristics render them incapable of delivering value which alternate approaches could achieve.
- A belief that deployment of PKI technology intrinsically implies and enforces a higher assurance environment than is appropriate or cost-effective in many operational contexts.

A 2003 survey undertaken by the OASIS PKI Technical Committee [Hann03] on obstacles to PKI deployment and usage suggests a mix of factors spanning each of these categories. If increased PKI adoption is taken as a goal, the first interpretation suggests a strategy of promoting applications and usage modes that would make use of certificates. Existing PKI technologies would stand ready to satisfy the demand if and as it emerges. While incremental changes might remain necessary to satisfy integration requirements, fundamental PKI architectures could safely remain intact. Questions of candidate applications and usages for PKI technology are interesting and important, but lie outside this paper's scope.

The second and third interpretations imply criticisms of elements within the PKI technology base, and motivations to revisit and modify those aspects of PKI that are considered to be contentious or problematic. Different commentators have expressed concerns about different elements of PKI technology, and have proposed different alternatives as a result; the goal of this paper is to examine a range of perceived issues and suggested approaches,

not to assert that all are equally valid or appropriate. Following this introduction, we characterize various perceived problem areas. Then, we examine several proposed approaches, seeking to characterize them in terms of the goals that they address, and the properties and value that they offer. We conclude by assessing asserted problems, and the contributions that suggested solutions make towards those problems.

This paper focuses on architectural and functional aspects of PKI. It is not primarily concerned with encoding alternatives, such as choices between ASN.1 and XML representations for protocol objects. For purposes of discussion, we assume the following elements as aspects of the contemporary PKI baseline, and therefore do not consider them under the category of candidate future variations:

- Support for hierarchic and non-hierarchic trust models
- Support for certificate revocation via Certificate Revocation Lists (CRLs) and via basic on-line status query mechanisms such as OCSP
- Syntactic support within certificates for a range of name forms, such as X.500 Distinguished Names, Internet-form names within AltName extensions, and pseudonyms.

While particular enhancements can be considered within many of these areas, their general premises have been widely presented and adopted, so do not constitute qualitative shifts from current accepted practice.

2 Contentious Aspects of PKI

In this section, we discuss several aspects of PKI technology and its operation that have attracted criticism and controversy.

2.1 Difficulty in Retrieving Keys and Certificates

To perform operations using public keys, those public keys must be available at the point where the operations are to be performed. In a conventional certificate-based PKI, this implies that a sender cannot encrypt a message for a recipient unless the recipient has already obtained a certificate and has made the certificate available to the sender (whether by direct transfer or posting on an accessible repository). If off-line operation is required, the appropriate certificates must be obtained in advance, when connectivity is available. Since large-scale directories have not become widely available to serve as certificate publication vehicles, interest has grown in approaches that enable public-key encryption to be performed without first satisfying these preconditions.

2.2 Questionable Value of Certified Key Representations

Certificates' usage practice reflects characteristics of environments for which they were originally developed, where it was considered inappropriate or impractical to rely on on-line availability of trusted servers. A primary goal of certificates' design was to represent keys and their bindings to named principals in an integrity-protected form, whose content could be stored safely on unprotected repositories or transferred across unprotected channels. Retrieval of a certificate requires that a suitable repository be available, but use of signed representations abstracts away the need to depend on that repository for security

properties other than availability. If, instead, keys are stored and retrieved from trusted servers, some of the rationale for representing them within signed certificate objects becomes superfluous. Channel-level mechanisms can protect a key from attackers while in transit between a server and a client, and can assure the client that it is receiving a key from a securely identified source.

2.3 Certificate Processing Complexity

PKI technologies have been criticized as being difficult to integrate with the applications that could make use of their services, requiring significant PKI-specific security expertise on the parts of application writers and maintainers. Today's X.509 certificates, e.g., have evolved into complex structures, with processing semantics that are far from trivial; this is primarily a matter of the information they carry, although it also involves its representation and encoding. Formalization and simplification of these semantics may represent a valuable area for investigation.

Some of the complexity in certification results from a desire for a certificate to include a comprehensive set of ancillary information so that it can be used for off-line processing, without consulting other trusted entities on an interactive basis. Increasingly, however, PKI models are evolving to include on-line components, which can offer alternative information sources to complement the certificates themselves.

Revocation mechanisms have long been recognized as a complex element in PKI, and path construction also introduces complexity [Elle01]. Despite the design attention that has been paid to revocation, it appears today that only a relatively small proportion of accepted certificates are actually checked for revocation status on an ongoing and timely basis.

2.4 Costly Certificates

Many assumptions about certificate usage have been based on a premise that certificates are expensive, and therefore that they can only be issued sparingly and infrequently. Some enrollment methods strive to provide confidence commensurate with high-value transactions and high-assurance client implementations, entailing high monetary costs and/or cumbersome registration processes. While this practice is appropriate for some types of technology (e.g., one-time placement of a user's long-term certificate into a smart card), and may be necessary to provide high levels of accountability, it need not be an intrinsic characteristic associated with the use of PKI methods. Imagine, by comparison, how computing might have developed if it had become accepted practice that an independent organizational authority needed to be consulted (and, possibly, paid) whenever a file was to be created. Most likely, only a subset of information, perhaps associated with a subset of critical users, would be deemed to warrant file representation. Other data would be stored and shared using different objects without the constraints associated with files. For a PKI, even when high levels of administrative assurance are not required, certification paradigms can be retained and adapted rather than developing or applying separate types of infrastructures to bind principals, keys, and attributes.

Dynamic issuance of certificates, which may be short-lived to avoid the need for separate revocation infrastructures, may allow new and innovative PKI models to be constructed.

In the Security Assertion Markup Language (SAML) [Male03], e.g., assertions bearing the Holder-of-Key confirmation method can take the form of signed objects carrying public keys, used to enable the corresponding private keys' holders to gain access to resources. Servers are expected to issue such assertions frequently, as needed to support authentication or resource access operations; no laborious procedures are required when an assertion is coined. Further, a number of on-line PKI key registration protocols (e.g., CMP [AdFa99], XKMS's X-KRSS [W3C03]) have been defined, which can provide the basis for interactive certification. The form of the resulting object, whether X.509, XML, or another format, need not imply or dictate the scope of procedural processing that is appropriate before the object is issued.

2.5 Problematic Cross-Domain Trust Management

The prospect of applying PKI technology to establish trust across heterogeneous domains can be daunting, both in administrative and technical terms. Some PKI architectures have sought to provide a sufficient basis to allow parties in different jurisdictions to engage in high-value transactions with one another, without prior shared knowledge beyond that manifested in the PKI. Few other technologies have attempted such ambitious goals, and it is debatable whether other approaches would necessarily achieve greater success in solving such a fundamentally challenging problem. In cases where the level of required assurance can be constrained, it may become easier to achieve (and benefit from) PKI-enabled interoperability.

PKI technologies can be applied to manifest trust relationships rooted at remote entities. Some (e.g., [DoE102]) have argued, however, that users' trust is primarily local, and should be based on direct personal knowledge of other individuals. If this premise is accepted, reliance on remote roots is not considered practical or useful, and the ability to represent such trust relationships offers only irrelevant complexity.

Meaningful algorithmic translation of policies across domain boundaries is a significant challenge; often, the mapping between different organizations' policy elements can be based on administrative practices and interpretations that are difficult to encode. Management of inter-domain validation and trust relationships within a relatively small set of entities (e.g., bridge CAs, domain-level Delegated Path Validation (DPV) servers interacting with their peers representing other domains) may help to contain and simplify some aspects of the problem.

2.6 Naming Semantics

Naming plays an important role in PKIs, as public keys are typically bound to named entities. Conventional PKIs have been criticized for seeking to manifest a global naming structure that some view as fundamentally unrealistic. As with trust, some view naming as intrinsically local; further, given duplications among human individuals' names, ambiguities can arise in identifying a particular person based on his or her location in a distributed namespace. In some alternate approaches, e.g., SDSI [RiLa96], entities are named in a relative manner extending from one principal, and then can be linked to other principals through intermediary hops.

Another aspect of PKI entity names is the degree to which a name form matches or resembles names that people and software use on a regular basis. This has a direct bearing

on how useful the name is to the user or application that is trying to accomplish a security goal. Some PKIs – such as Pretty Good Privacy (PGP) [Call98] and the DNS security extensions (DNSSEC) [East99] – employ name forms that match their environments (or, rather, they adopt the name form of their environment). X.509 is an example of a PKI that started out adopting the name form of its environment (X.500 Distinguished Names), but then grew to accommodate application-specific names (through the Alternative Name extensions). A SDSI “well-defined” name – one that links a local name space to a particular principal, such as (using SDSI’s “syntactic sugar”) `jim's john's joe's jack` – is only meaningful to the SDSI PKI. However, each individual local name is an arbitrary string, and so can be meaningful to an application. For example, `10.1.1.1` might be a local SDSI name assigned to a VPN server whose IP address is, presumably, `10.1.1.1`. PKIs with PKI-specialized name forms require applications to translate between their native name form and the PKI's, a process that can be error-prone and introduce security risks.

PKI	Name Locality	Name Form Application	Utility PKI
PGP	Low	High	Low
DNSSEC	Low	High	High
X.509	Low	High	Low
SPKI/SDSI	High	Medium	High

Table 1 - PKI naming properties

A third property of PKI names is the degree of utility that the name has to the PKI itself. By "degree of utility" we mean the efficiency with which the PKI can use the name to obtain and validate a public key. PKIs provide keys by discovering and validating paths between entities, and so the PKI-efficiency of a name can be measured by the amount of path information that it encodes. Well-defined SDSI names are an extreme example of a name form that is almost entirely devoted to expressing path information, so much so that a (non-global) SDSI name is usually only meaningful to a single entity. DNSSEC names also encode a large amount of path information. In contrast, PGP names are email addresses, which are completely devoid of any PGP PKI path data. X.509's names – all of them – also contain no X.509 path information whatsoever.

Table 1 summarizes the naming properties for various PKIs. SDSI scores highly for name form PKI utility because of its well-defined names, but only moderately for application utility because although an individual local name can be an application-meaningful string, there are no conventions for an application to reliably extract a meaningful local name from a SDSI certificate. A VPN client, for example, has no way to tell that the `10.1.1.1` name in a SDSI certificate is supposed to be the IP address of a VPN server.

Recent PKI proposals have emphasized certificate processing and cryptographic methods rather than naming. A viable naming strategy seems to be a factor in a PKI's success, but it is not clear what combinations of properties (per Table 1) offer most value. Naming strategies do appear to require some consideration, and yet they remain relatively unexplored. Some of the questions that arise include:

- Are there any other useful naming properties?
- Is it necessary or desirable to rank highly in all of these properties?
- Have approaches to naming had an impact on PKI deployment? We note, for example, that an X.509 certificate in fact has two names – Issuer and Subject

- which together provide a small amount of path information. Would more (or less) path information in the certificate help or hinder widespread deployment of an X.509 PKI?

2.7 Use with Insecure Clients

Some PKI architecture premises were developed in anticipation of widespread security features at user clients, e.g., smart cards encapsulating users' private keys and cryptographic processing capabilities so that the keys need never be exposed elsewhere. Such implementations are particularly desirable when the keys mediate access to particularly sensitive data or resources, or when strong accountability (i.e., a non-repudiation service) is tied to their use. While such environments are gradually becoming more common (as with use of SIMs and other cards), most candidate PKI user applications continue to reside on platforms that offer limited security. From an attacker's viewpoint, the strength of a cryptographic algorithm can become irrelevant if its keys can be obtained by attacking a weak platform. Where high assurance is required, these arguments motivate approaches that perform cryptographic processing in other entities, whether protected devices or shared services, and/or distribute the processing with such entities.

There are many cases, however, where the assurance level of commercial platforms is an adequate basis to support useful, interoperable security. Use of PKI need not also imply use of specialized, higher-security technologies by clients; higher assurance requirements may be warranted at CAs, as misuse of a single CA private key can compromise an entire community. Today, it is common practice to store user keys in a password-encrypted form. It is arguable that passwords used to unlock private keys may warrant higher quality or tighter protection than other passwords, as the keys they release can enable direct authentication to multiple entities rather than just to a single system, but user convenience may conflict with such measures.

2.8 Privacy Compromises

It has been observed, e.g., in [Bran99], that conventional PKI is unfriendly to privacy, as its certificates provide persistent, widely visible linkages between keys and principal identifiers. This property is appropriate in contexts where authorizations or signatures depend on individuals' authenticated identities, but not all possible uses of public-key technology fit this model. Even if data messages are encrypted, patterns of certificate acquisition and usage can reveal identities of principals and their communicating peers; a certificate validation server could be particularly well placed to collect such information. Certified pseudonyms can provide a partial countermeasure, but do not satisfy all privacy goals; if a fixed pseudonym is used to represent a principal to multiple sites for an extended period, the sites can use it as the basis to collect an extensive behavior profile which may then be associated with an individual.

Use of X.509 certificates to hold principal attributes other than identifiers has been proposed and considered for some time, recently in [FaHo02], though has not yet achieved wide adoption. Attribute statements within SAML assertions are another form of attribute representation within a signed object corresponding to a principal. Both have the property of disclosing an aggregate set of attributes to their certifier and to the parties that rely

on the certified object, even if not all of these entities necessarily require the full set of information.

3 Proposed Approaches

In this section, we examine approaches that have been proposed as extensions or alternatives to conventional PKI technologies, addressing one or more of the concerns identified in the preceding section.

3.1 IBE and Related Work

The concept of Identity-Based Encryption (IBE) has been considered in the cryptographic community for some time, and recent work has yielded a variety of methods realizing variations on the concept. Some, but not all, approaches in this group allow a sender to prepare a protected message for a recipient without first obtaining a certificate for the recipient. This section considers some of their properties.

3.1.1 Identity-Based Encryption

IBE, surveyed in [Gagn03], enables senders to encrypt messages for recipients without requiring that a recipient's key first be established, certified, and published. The basic IBE paradigm allows a sender to determine the key to be used to encrypt for a particular recipient based on the recipient's identifier; the recipient derives the corresponding decryption key through interaction with a Private Key Generator (PKG) system. While the sender must determine the PKG corresponding to a particular recipient, and must obtain domain-level parameters associated with that PKG, it need not obtain information specific to an individual recipient before encrypting a message. The basic IBE approach implies intrinsic key escrow, as the PKG can decrypt on behalf of the user. Variant approaches ([AlPa03] [Gent03]) cited below apply some aspects of IBE, but seek to avoid the escrow characteristic.

3.1.2 Certificateless Public Key Cryptography

This approach, proposed in [AlPa03], incorporates IBE methods, using partial private keys so the PKG can't decrypt on behalf of the user. These are combined with secret information held by the recipient, yielding a public key that the recipient can publish and/or transfer directly, but for which no certification is required. Would-be senders must, however, first obtain that key through some means in order to encrypt a message for a recipient. Publication of a key for this method may not prove significantly easier than publishing a conventional PKI certificate. In fact, the publication problem could become significantly worse, since use of the approach might imply a need for frequent republication in lieu of a revocation mechanism.

3.1.3 Certificate-Based Encryption

This approach, proposed in [Gent03], incorporates IBE methods, but uses double encryption so that its CA can't decrypt on behalf of the user. A sender must obtain a recipient's certificate in order to encrypt a message for a recipient. In order for a recipient to decrypt successfully, he/she must have both a current CA-issued certificate and a personal secret key; use of IBE methods in certificate generation means that the same certificate used by

the sender to encrypt is also used by the recipient as part of the decryption process. Frequent certificate updates are performed, so that senders need not separately check revocation status of the certificates they obtain.

3.2 PKI Augmented with On-Line TTP

Some properties similar to those of IBE can be achieved by augmenting conventional PKI with an on-line trusted third party (TTP) system. Two classes of TTP-based operations can be considered:

- Encryption using a TTP's public key rather than one associated with an individual recipient; in this case, a recipient could request that the TTP perform decryption services on his/her behalf, or a message could be routed to the TTP which would then decrypt it and forward the result to the recipient. This eliminates the need for recipients to register individual key pairs, and for senders to obtain per-recipient keys; it implies that the TTP can decrypt all recipients' traffic and requires involvement by the TTP in order to process each of their messages. [DeOt01] provides examples and discussion of this type of approach.
- Encryption using an individual recipient's public key, which the sender would request from the TTP. For already-registered recipients, a TTP (such as that suggested in [Dier03]) would provide their existing keys or certificates. Additionally, such a TTP could revoke keys or certificates by removing them from its store. If no public key or certificate existed for the recipient at the time of the request, the TTP would generate one dynamically, provide the public component to its requester, and make the corresponding private key available to the recipient. In this model, the TTP's possession of recipients' private keys need not be more than temporary in nature, pending their retrieval by the corresponding recipient.

The second type can be considered as an example of a general class which has previously been considered in various contexts but has not become part of the PKI mainstream, that of "on-the-fly PKI" approaches where certificates are signed dynamically as needed rather than being generated by a CA in advance as a prerequisite to secure operation. Such certificates and the keys they certify can be short-lived, enabling particular operations or use of a session while becoming disposable thereafter. Some other examples include the delegation certificates that represent login sessions within Digital Equipment Corporation's Distributed System Security Architecture (DSSA) as proposed ca. 1990 [GaMcD90], and recent IETF-PKIX contributions on proxy certificates [Tuec03].

3.3 Distributed Computation

Methods have been developed (see, e.g., [Gold02]) that distribute cryptographic operations so that the cooperative contribution of a number of entities is required in order to perform an operation such as a signature or a decryption. Use of such measures could help to ameliorate the risks associated with insecure client platforms; even if such a client's keys were compromised, they would be insufficient to impersonate the client's associated user.

Analogous to the case with IBE, some similar properties can also be achieved without specialized cryptography by holding a user's keys at a server, which would perform operations on behalf of the user upon receipt of an authenticated request. This strategy can take advantage of tighter protection at servers vs. clients, but implies that the users must fully trust the servers to apply their keys appropriately.

3.4 Alternative Validation Strategies

PKI's original Certificate Revocation List (CRL) mechanisms implied significant storage, communications bandwidth, and processing overhead, yet could only provide revocation with significant time latency. Newer on-line approaches, such as OCSP, SCVP, and XKMS, address many of these concerns, but introduce requirements for trusted on-line servers to process certificates and for connectivity between the servers and their relying parties. Their effective revocation latencies can vary, as a result of caching and when information updates are available only on a periodic basis. These approaches' capabilities, and the extent to which clients must trust the servers, increase as the scope of server-based processing extends from revocation checking on single certificates to acquisition and validation of full certification paths, and from independent, self-contained validation servers to distributed networks of cooperating validators. More broadly, however, the extent of trust required should correspond to the value of the information that the underlying certificates protect. Further discussion of validation alternatives and their prospects and implications can be found in [BrLi02].

Hash-tree approaches (e.g., [Mica02] [NaNi98]) have been proposed, offering compact, protected representations of the status of large numbers of certificates. Their value is most apparent for PKIs operating at extremely large scale; in smaller contexts, such as within typical enterprises, their benefits relative to CRLs appear less compelling. Like CRLs, they reflect certificate status information only at fixed intervals, rather than with the immediacy that on-line status queries can offer.

Levi and Caglayan [LeCa00] propose the concept of "nested certificates" in order to avoid some of the performance burdens associated with verification of long certification paths. Several variations are suggested, but a general premise is that a hierarchy's higher-level CAs certify not only their immediate descendants but also directly certify members of more distant generations. While this approach can indeed reduce the number of certificates in a validated path, it appears to suffer from a serious flaw. Among other reasons, CA hierarchies are constructed in order to distribute certification responsibilities, and to place them in hands close to the principals they certify. In a condensed hierarchy, higher-level CAs would need to be involved in enrollment of remote generations, and potentially to generate very large numbers of certificates. In the limit, a CA hierarchy could be flattened to a single CA, making any hierarchy below it moot, but such an approach is unlikely to be attractive from a technical or policy perspective.

3.5 Key Servers

Given today's generally high level of connectivity, and widespread interest in simplifying client-side operations, an emerging approach is to use servers to perform some, or all, certificate processing. Clients would delegate certificate path discovery and/or validation to a trusted server (see [PiHo02]).

DPV, in particular, changes the basic PKI model. A DPV server assumes the primary responsibilities of a traditional CA, from the client's perspective. That is, the client relies upon the server to ensure correct correspondences between principals and their public keys.

This approach has implications for assurance and availability, especially when a DPV server relies on other DPV servers (see [BrLi02]). However, once the premise of trusting an on-line server for certificate retrieval and validation is accepted, it is only an incremental step to relying on the server to provide the bare key over a secure channel – eliminating the need for the client to process certificate formats entirely. Such an approach is one of the models supported within XKMS's X-KISS [W3C03].

The full impact of delegating key validation and acquisition to servers has yet to be investigated. The benefits to PKI client applications for smaller, simpler code are apparent, but it is not yet clear what effects delegated key servers will have on a PKI's policies and procedures, or what levels of assurance are enabled (or disabled).

3.6 Privacy Protection

Some PKI privacy implications can be ameliorated by reducing the amount of principal-related information bound within a single certificate or other signed object. Certified pseudonyms can easily be supported, and are appropriate and sufficient in many operational contexts. Further, use of attribute certificates (ACs) can offer privacy advantages over placement of attributes within public-key identity certificates (PKCs). Even in the common case where an AC is bound to a PKC for use, implying a linkage to the PKC within the AC, the PKC's contents need not disclose all of the ACs that may be used with it. This modularity allows the attributes within ACs to be disclosed selectively, when needed in order to support a particular access request, and to remain confidential otherwise. To take advantage of this capability, it is desirable for accessors to present ACs selectively along with requests rather than posting them for general access within a directory or other repository.

Use of on-line certificate validation services introduces the prospect of user tracking, if the validation service can identify the set of locations from which a certificate's status is queried. Aggregation and/or anonymization of status requests can help to mitigate this concern.

Stefan Brands, in [Bran99], proposes cryptographic certification techniques which address privacy goals outside the scope of traditional PKI models, and which imply different assumptions and paradigms for PKI protocols and interactions. Brands' techniques seek to allow certificate holders to disclose certified attributes selectively in a general manner, and to limit the extent to which presentation of certified attributes can be proven to third parties by recipients. Cryptographic blinding is used for certificate issuance, so that not all of the attributes represented within a certificate need be visible to a particular issuing CA. These approaches can provide privacy assurance unavailable in conventional PKIs, particularly in terms of constraining the scope of trust that a certified user must place in a CA and of countering use of certificates as a means to aggregate data. Their operational models would require changes in certificate-based protocols, one factor which would likely complicate their deployment.

4 Conclusions

Any concrete system can suffer in comparison with a hypothetical, ideal alternative. PKI has been a particularly attractive target, perhaps partly because it has sometimes been perceived and promoted as a general panacea, intended to solve even organizational issues outside the realm of technology, rather than as a technical answer to clearly understood and practically achievable requirements. Variations to many aspects of PKI are possible and worthy of consideration, but an appropriate comparison between practice and proposal requires a specific alternative and an understanding of its impact on the system as a whole.

Certificates have been criticized for a variety of reasons, particularly:

- Processing complexity and overhead, including both the contents of certificates and the usage of signed representations to carry those contents; many of these characteristics derive from design assumptions which presumed off-line certificate processing without reliance on trusted servers, and use of such servers may allow significant simplifications.
- Association with operational models that imply high costs for certificate issuance; here, the use of a signed key-bearing object should properly be distinguished from a particular type of deployment. Public-key methods can be used to construct a wide variety of useful approaches with different assurance, semantics, and dynamics.

PKI has also been criticized on the basis that it fails to render the problems of securely interconnecting different entities and trust domains simple. These problems are fundamentally difficult, for organizational as well as technical reasons. Few proposals outside the realm of PKI have attempted to satisfy these concerns comprehensively, though trust management research activities [Blaz99] have proposed various supporting mechanisms. Generally, PKIs' trust management capabilities should be evaluated in terms of their supporting contributions to distributed security, rather than against an expectation that all such requirements should be satisfied solely by PKI or any other technology.

Much cryptographic research activity has concerned forms of IBE, applied to avoid the need for senders to retrieve certificates from repositories. Unfortunately, many proposed alternatives substitute different publication requirements, or introduce implicit key escrow properties. Other computational methods can distribute processing, mitigating some of the impact of key compromise at weakly protected clients. These cryptographic innovations provide elegant approaches, but many of their properties can also be achieved by using trusted third parties with conventional cryptographic algorithms.

Fundamentally, PKIs exist to provide public keys that correspond to principals, in a fashion enabling other parties to rely on their correspondence. This function is an essential basis on which to construct secure distributed computing environments, and necessarily implies some form of infrastructure. Many PKIs seek to provide high levels of technical and procedural assurance, particularly at CAs, but some of these measures may not be necessary for environments where the ability to communicate with at least some level of protection takes precedence over especially strong security guarantees. Naming is a cen-

tral element in PKI, and further research focused on aspects of alternate naming methods may warrant attention.

Certificates are a convenient, self-sufficient means of representing keys, but their use may become superfluous in server-centered environments. Further, new PKI models can evolve based on signed key-bearing assertions; these objects can provide the same functions as certificates, but are emerging unbounded by existing assumptions about how certificates must be created, processed, and managed. Generally, it seems that PKI suffers today from a perception that it can assume only a particular, monolithic form; to satisfy a broad range of applications and environments, it must be possible for its underlying methods to be composed and applied in a variety of ways.

5 Acknowledgment

The authors would like to acknowledge this paper's anonymous reviewers for comments helping to improve its final version.

6 References

- [AdFa99] C. Adams, S. Farrell, "Internet X.509 Public Key Infrastructure Certificate Management Protocols", Internet RFC-2510, March 1999.
- [AlPa03] S. S. Al-Riyami and K. G. Paterson, "Certificateless Public Key Cryptography", IACR Cryptology ePrint Archive paper 2003/126, 2 July 2003.
- [Blaz99] M. Blaze, J. Feigenbaum, J. Ioannidis, A. D. Keromytis, "The Role of Trust Management in Distributed System Security", in *Secure Internet Programming: Issues in Distributed and Mobile Object Systems*, Springer-Verlag Lecture Notes in Computer Science State-of-the-Art Series, pp. 185-210, Berlin, 1999.
- [Bran99] S. Brands, "Rethinking Public Key Infrastructures and Digital Certificates – Building in Privacy", PhD Dissertation, University of Utrecht, October 1999.
- [BrLi02] M. Branchaud, J. Linn, "Extended Validation Models in PKI: Alternatives and Implications", 1st PKI Research Workshop, Gaithersburg, MD, April 2002.
- [Call98] J. Callas, et al., "OpenPGP Message Format", Internet RFC-2440, November 1998.
- [Clar01] R. Clarke, "The Fundamental Inadequacies of Conventional Public Key Infrastructure", Proceedings, ECIS'2001, Bled, Slovenia, June 2001. Available at <http://www.anu.edu.au/people/Roger.Clarke/II/ECIS2001.html>. (Date of access: 26 November 2003.)
- [DeOt01] T. Dean, W. Ottaway, "Domain Security Services Using S/MIME", Internet RFC-3183, October 2001.
- [Dier03] T. Dierks, "Re: Fwd: [IP] A Simpler, More Personal Key to Protect Online Messages". Message posted to Cryptography electronic mailing list, ar-

chived at <http://www.mail-archive.com/cryptography@metzdowd.com/msg00409.html>. (Date of access: 4 December 2003.)

- [DoEl02] S. Dohrmann, C. Ellison, “Public-key Support for Collaborative Groups”, 1st PKI Research Workshop, Gaithersburg, MD, April 2002.
- [East99] D. Eastlake, “Domain Name System Security Extensions”, Internet RFC-2535, March 1999.
- [Elle01] Y. Elley, et al., “Building Certification Paths: Forward vs. Reverse”, NDSS-01, San Diego, 2001.
- [ElSc00] C. Ellison, B. Schneier, “Ten Risks of PKI: What You’re Not Being Told about Public Key Infrastructure”, Computer Security Journal, Vol. XVI, No. 1, 2000. Available at <http://www.schneier.com/paper-pki.html>. (Date of access: 4 March 2004.)
- [FaHo02] S. Farrell, R. Housley, “An Internet Attribute Certificate Profile for Authorization”, Internet RFC-3281, April 2002.
- [GaMcD90] M. Gasser, E. McDermott, “An Architecture for Practical Delegation in a Distributed System”, Proceedings, IEEE Computer Society Symposium on Research in Security and Privacy, Oakland, CA, May 1990.
- [Gagn03] M. Gagné, “Identity-Based Encryption: a Survey”, RSA Laboratories Cryptobytes, Vol. 6, No. 1, Spring 2003.
- [Gent03] C. Gentry, “Certificate-Based Encryption and the Certificate Revocation Problem”, EUROCRYPT 2003, LNCS 2656, pp. 272-293, 2003.
- [Gold02] O. Goldreich, “Secure Multi-Party Computation (Final (Incomplete) Draft Version 1.4)”, 27 October 2002. Available at <http://www.wisdom.weizmann.ac.il/~oded/pp.html>. (Date of access: 19 December 2003.)
- [Gutt02] P. Guttman, “PKI: It’s Not Dead, Just Resting”. Available at <http://www.cs.auckland.ac.nz/~pgut001/pubs/notdead.pdf>. (Date of access: 5 March 2004.)
- [Hann03] S. Hanna, ed., “Analysis of August 2003 Follow-up Survey on Obstacles to PKI Deployment and Usage”, OASIS PKI Technical Committee, 1 October 2003. Available at <http://www.oasis-open.org/committees/pki/pkiobstaclesaugust2003surveyreport.pdf>. (Date of access: 4 March 2004.)
- [LeCa00] A. Levi, M. Caglayan, “An Efficient, Dynamic, and Trust Preserving Public Key Infrastructure”, Proceedings, IEEE Computer Society Symposium on Research in Security and Privacy 2000. IEEE, Piscataway, NJ, USA, pp. 203-214.
- [Male03] E. Maler, P. Mishra, R. Philpott, eds. (2003), “Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V1.1.” OASIS Standard.

- [Mica02] S. Micali, “Novomodo: Scalable Certificate Validation and Simplified PKI Management”, 1st PKI Research Workshop, Gaithersburg, MD, April 2002.
- [NaNi98] M. Naor, K. Nissim, “Certificate Revocation and Certificate Update”, 8th USENIX Security Symposium, San Antonio, January 1998.
- [PiHo02] D. Pinkas, R. Housley, “Delegated Path Validation and Delegated Path Discovery Protocol Requirements”, Internet RFC-3379, September 2002.
- [RiLa96] R. Rivest, B. Lampson, “SDSI – A Simple Distributed Security Infrastructure”, 30 April 1996.
- [Tuec03] S. Tuecke, V. Welch, D. Engert, L. Pearlman, M. Thompson, “Internet X.509 Public Key Infrastructure Proxy Certificate Profile”, work in progress, IETF PKIX working group, 2003.
- [W3C03] World Wide Web Consortium, “XML Key Management Specification (XKMS)”, Version 2.0, W3C Working Draft, 18 April 2003.

An Examination of Asserted PKI Issues and Proposed Alternatives

John Linn (RSA Laboratories), Marc
Branchaud (RSA Security)

3rd PKI R&D Workshop

NIST

April 2004

DRAFT, 4 April 2004

Presentation Structure

- Introduction
- Asserted Issues
- Proposed Approaches
- Conclusions

Motivations

- Many issues have been raised about PKI over time
- Many variations have been proposed
- Which issues are valid?
- Do the alternatives resolve the issues?
- How would useful results appear in context?

Why PKI isn't pervasive: three possible causes

- PKI technology is suitable, but awaits motivating demand for secure applications
 - Important hypothesis to consider, but not this paper's focus
- PKI technologies are hard to deploy, or deliver limited or less-needed value
- PKI is perceived to imply and require higher assurance than is necessary in many environments

Some Elements Assumed in Contemporary PKI Baseline

- Support for hierarchic and non-hierarchic trust models
- Revocation via CRLs or basic on-line queries like OCSP
- Support for various name forms
- NB: Once upon a time, each of these were novel...

Issue: Difficult to Retrieve Keys and Certificates

- It can be hard to obtain a PKI certificate
 - If no directory exists for publication
 - If that directory isn't accessible or can't be located
 - If the certificate is needed by an off-line user
- Interest in avoiding or simplifying public-key processing preconditions

Issue: Value of Certified Keys

- The trusted key is the goal, and certificates are a mechanism
- Certificates were developed to represent keys in a protected fashion on an untrusted repository
- If keys are obtained over a secure channel from an on-line trusted server, value of certification diminishes

Issue: Certificate Processing Complexity

- It's complex to process certificates, and to integrate their processing with applications
 - Interpreting key usage indicators, ...
- Validation and path-level processing add further complexity
- Have certificates grown to include an unwieldy amount of free-standing data?

Issue: Costly Certificates

- Common assumption: certificates are expensive, so can only be issued rarely and sparingly
- High-assurance enrollment procedures are appropriate in some contexts, but not always needed
- Certificates/assertions can be issued dynamically

Issue: Problematic Cross-Domain Trust Management

- Enabling trust between unrelated entities is a daunting challenge, administratively and technically
- Conventional PKI reflects trust between domains, not between principals
- Policy mapping provides mechanism, but may not fit practices
- Manual trust anchor management is a common and limiting constraint

Issue: Naming Semantics

- PKIs bind keys to names, but not all names have the same properties
 - PKI names can be local vs. global
 - PKI names can match or diverge from names used in other contexts
 - PKI names can imply paths or can be independent of them
- What properties are most useful?

Issue: Use with Insecure Clients

- PKI designs have anticipated deployments where users securely control keys
- Many common platforms are subject to compromise
- Can provide useful security even in “commercial practice” environment
 - Various methods can improve assurance
 - Even without perfect client protection, PKI-based services can still be useful

Issue: Privacy Compromises

- Conventional PKI certificates provide signed linkages between principals and actions
- Persistent keys become identifiers
- Certified identities can be used for profiling
 - Validation servers can be well-placed observers...

Issues to Proposals

- Several categories of issues have been asserted
- Several types of proposals have been made, responding to different concerns
- Goal: consider value and implications

Proposal: IBE and Related Methods

- Several variants on Identity-Based Encryption have been defined
- Many allow a sender to prepare a message for a recipient without obtaining the recipient's certificate
- Sender needs parameters for recipient's domain, implying need for cross-domain infrastructure
- Basic IBE approach implies intrinsic key escrow

Proposal: PKI with On-Line TTP

- On-line TTPs can achieve IBE-like properties
 - Encrypting with a TTP's public key
 - Encrypting with a recipient's key, which the TTP can provide (or generate)
- Dynamic certificate generation can also serve other purposes (temporary attributes, login sessions)

Proposal: Distributed Computation Methods

- When platform compromise or constrained trust is an issue, can limit impact
 - By storing principal keys on a protected server, requesting remote operations
 - By distributing key elements and performing cooperative computation
- Provide assurance at overall system level, rather than per-component

Proposal: Alternative Validation Strategies

- CRLs operate off-line, but provide coarse revocation latency
- On-line services can provide finer latency, but require trust and availability
- Hash trees optimize CRL-like properties, with particular value at large scale

Proposal: Key Servers

- If on-line servers are fully trusted for path-level discovery and validation of certificates, it's an incremental step for them to provide keys directly
- Clearly simplifies clients, but also changes assurance model and assumptions

Proposal: Privacy Protection Approaches

- Can certify (temporary) pseudonyms
- Can separate attributes into individual certificates, presented selectively
- Can aggregate or anonymize status queries
- Alternate certification models provide qualitatively stronger protection, but could require new operational paradigms

Observations

- PKIs have addressed broad and difficult problems with partial success
 - Technologies can reflect organizational conflicts, can't generally resolve them
- Self-contained certificates allow off-line processing, but with management and complexity costs
 - Useful to decouple assumptions about certificate properties

Conclusions

- PKIs, in suitable forms, remain essential substrates for secure transactions
 - Need various means to securely provide keys for different contexts
- Methods and their assurance levels should reflect requirements, need not lead them
 - Conventional vs. dynamic vs. no certificates...
 - Appropriate tradeoffs among client protection, client processing, server trust

Which PKI Approach for Which Application Domain?

3rd Annual PKI R&D Workshop

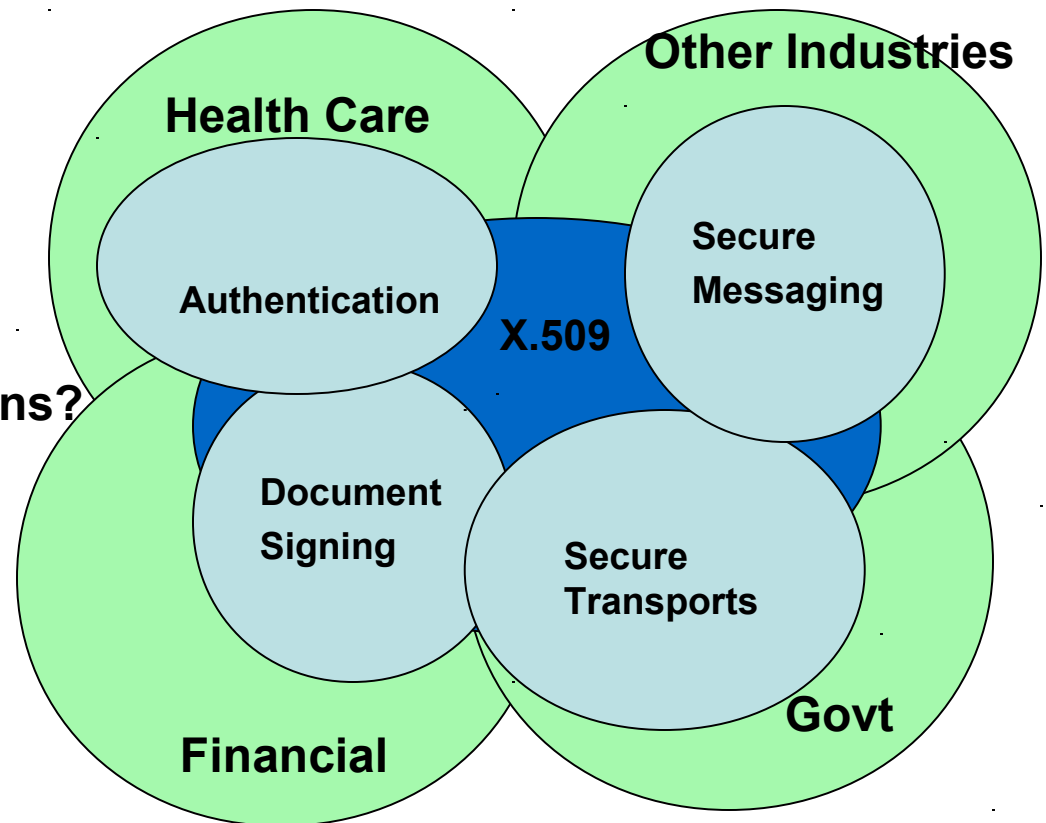
Russel F Weiser

April 2004

Which PKI? -- X.509 A Perspective

Which Applications?

- Document Signing
- Authentication
- Secure Messaging
- Secure Transports



What Industries or Trust Domains?

- Government
- Health Care
- Financial
- Others

- **Low Risk documents** - Time Cards / Standard Reports
- **High Risk documents** – Contracts / Financial Transactions
- **Complex documents** - HR Performance Reviews
- **Document Management** – Work Flow / Document Control

Spans all of Industry Domains

Example – Access to patient record information for research including “Patient Release Authorization”

Application	Govt.	Health Care	Financial	Other Industries
Document Signing				
Low Value		X		
High Value	X	X	X	X
Complex (Workflow)	X	X	X	X

- **S/MIME**
- **Instant Messaging – IMing**

Email touches all Industry Domains, Instant Messaging is getting there?

Example – The Financial Industry in particular is leveraging “Secure IM” across multiple enterprises.

Application	Govt.	Health Care	Financial	Other Industries
Secure Messaging				
Email	X	X	X	X
Instant Messaging (IM)		X	X	X

- **Network Authentication – Signal Sign On etc**

Spans all of the industry domains

Examples - Authenticating to corporate and partner networks is and obvious area where trust is across the enterprises.

Note – Several areas that would certainly benefit from any of the PKI solutions

- **Online Banking Access**
- **Retirement Management Access**
- **Brokerage Accounts**

Application	Govt.	Health Care	Financial	Other Industries
Authentication				
Network authentication	X	X	X	X

- **Web server Security - SSL/ TLS**
- **VPNs and IPSEC**
- **Secured EDI - Over HTTPS or SFTP**

Spans all of the Industry Domains

Example – Secure transfer of medical information via HTTPS based EDI.

Application	Govt.	Health Care	Financial	Other Industries
Transport Security				
Web Server Security	X	X	X	X
Virtual Private Network	X	X	X	X
EDI	X	X	X	X

Questions



Contact Information

Russel F Weiser
Betrusted US Inc.
Managing Consultant
Rweiser@betrusted.com
Office 801-942-6480
Cell 801-631-1685



A New and Improved Unified Field Theory of Trust

Ken Klingenstein,
Director, Internet2 Middleware and Security

- PKI Conference 03: REthinking Trust Flashback
- New and improved unified field theory of trust
- Looking at the data
 - Federation
 - P2P
 - Virtual organizations and authorizations
- Next year's talk...
 - Interfederation Issues
 - Federated PKI
 - Progressive PKI
 - Diagnostic hell

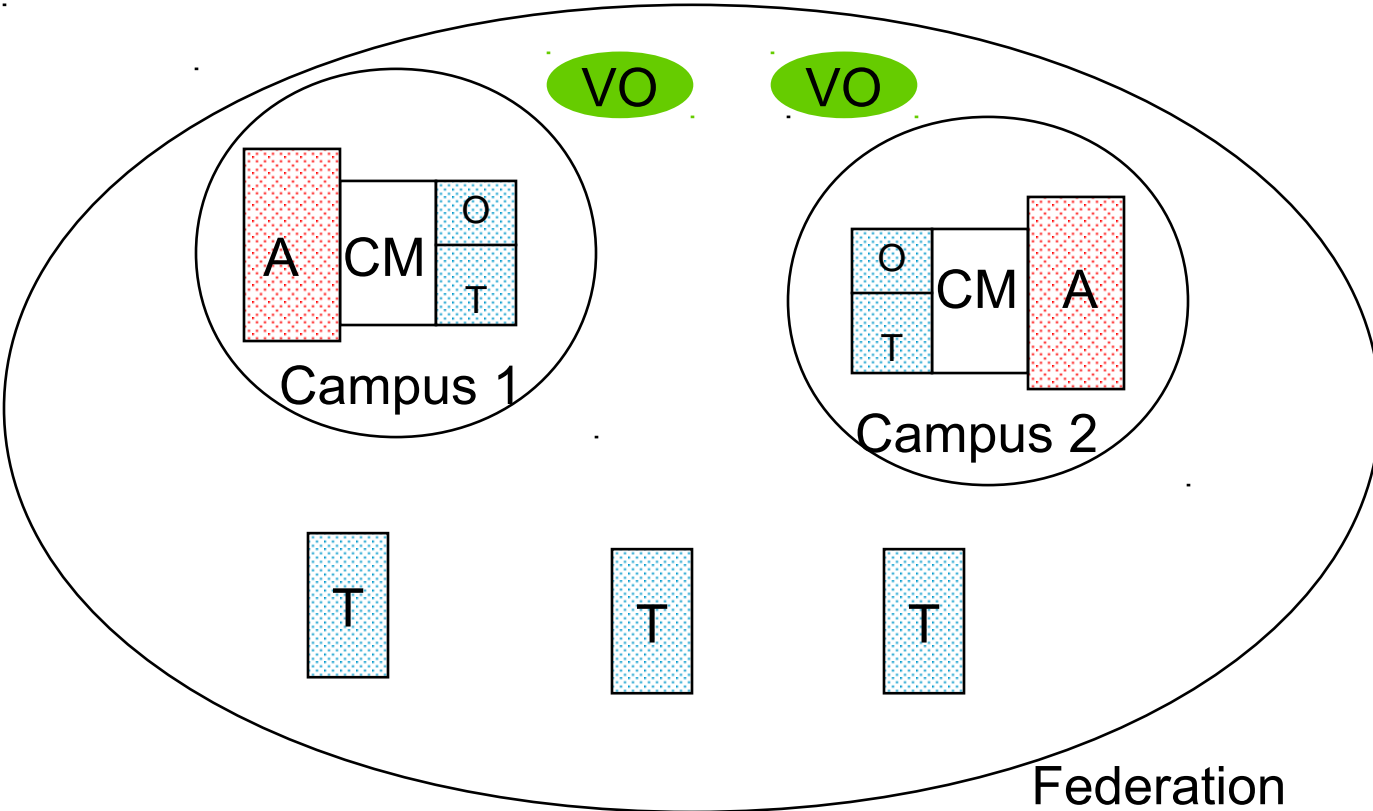
Unified field theory of Trust

- Bridged, global hierarchies of identification-oriented, often government based trust – laws, identity tokens, etc.
 - Passports, drivers licenses
 - Future is typically PKI oriented
- Federated enterprise-based; leverages one's security domain; often role-based
 - Enterprise does authentication and attributes
 - Federations of enterprises exchange assertions (identity and attributes)
- Peer to peer trust; ad hoc, small locus personal trust
 - A large part of our non-networked lives
 - New technology approaches to bring this into the electronic world.
 - Distinguishing P2P apps arch from P2P trust
- Virtual organizations cross-stitch across one of the above

Federated administration

- Given the strong collaborations within the academic community, there is an urgent need to create inter-realm tools, so
- Build consistent campus middleware infrastructure deployments, with outward facing objectclasses, service points, etc. and then
- Federate (multilateral) those enterprise deployments with interrealm attribute transports, trust services, etc. and then
- Leverage that federation to enable a variety of applications from network authentication to instant messaging, from video to web services, from p2p to virtual organizations, etc. while we
- Be cautious about the limits of federations and look for alternative fabrics where appropriate.

Federated administration



Federation

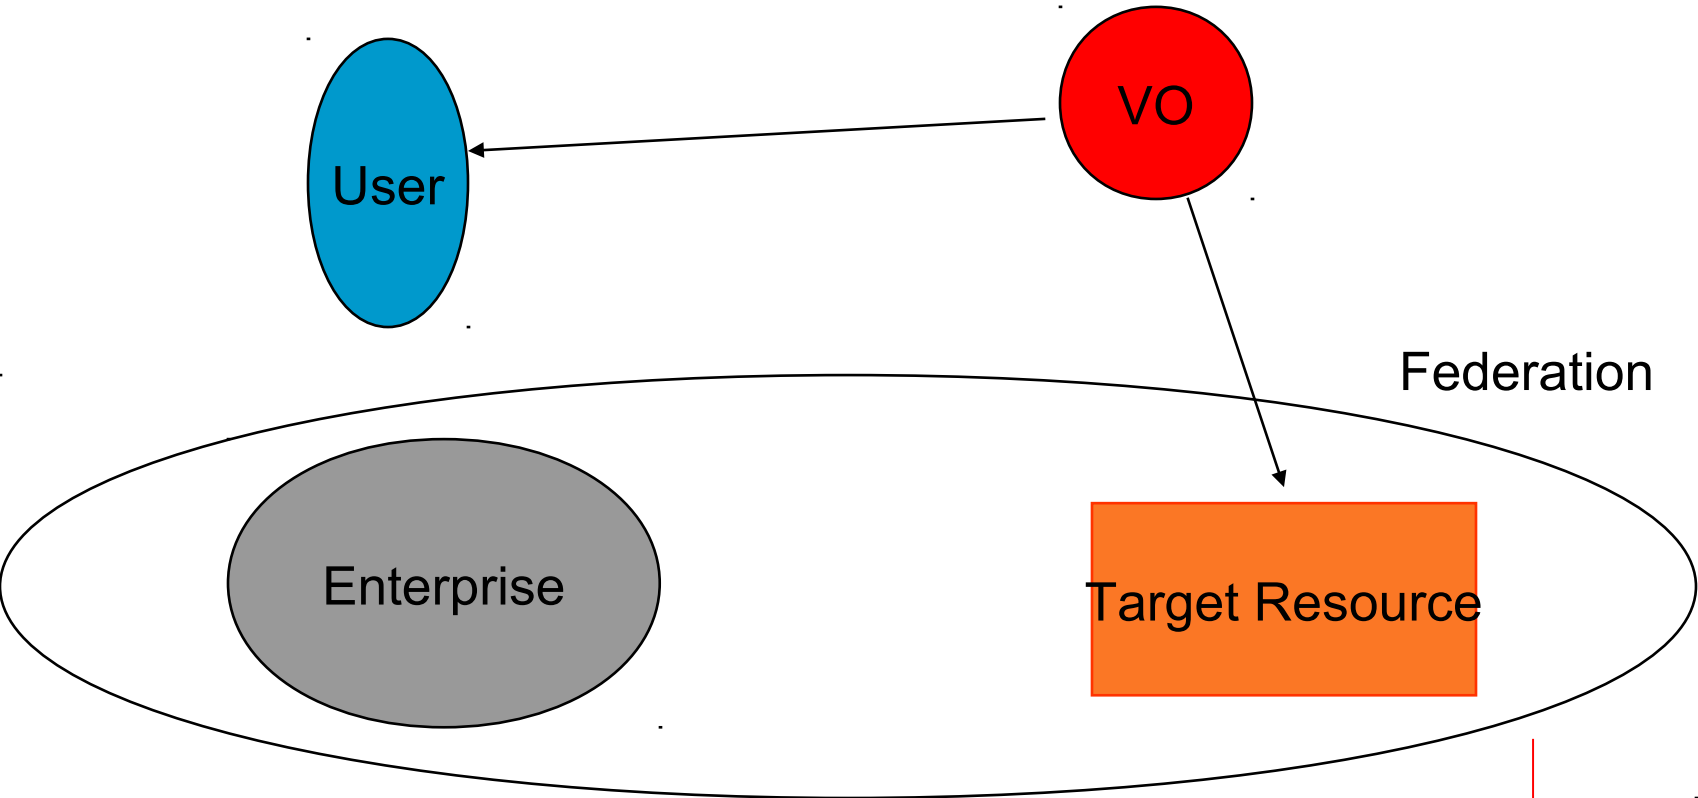
Peer to peer trust

- A bedrock of human existence
- Completely intuitive, sometimes contradictory and soft around the edges
- Translation into technology is difficult
 - PGP and webs of trust most successful
 - X.509 Proxy Certs a new, odd option
 - Issues over transitivity, integration into applications, user management are hard
- Some new technologies, embedded within MS Longhorn, present an option that will have a large embedded base...

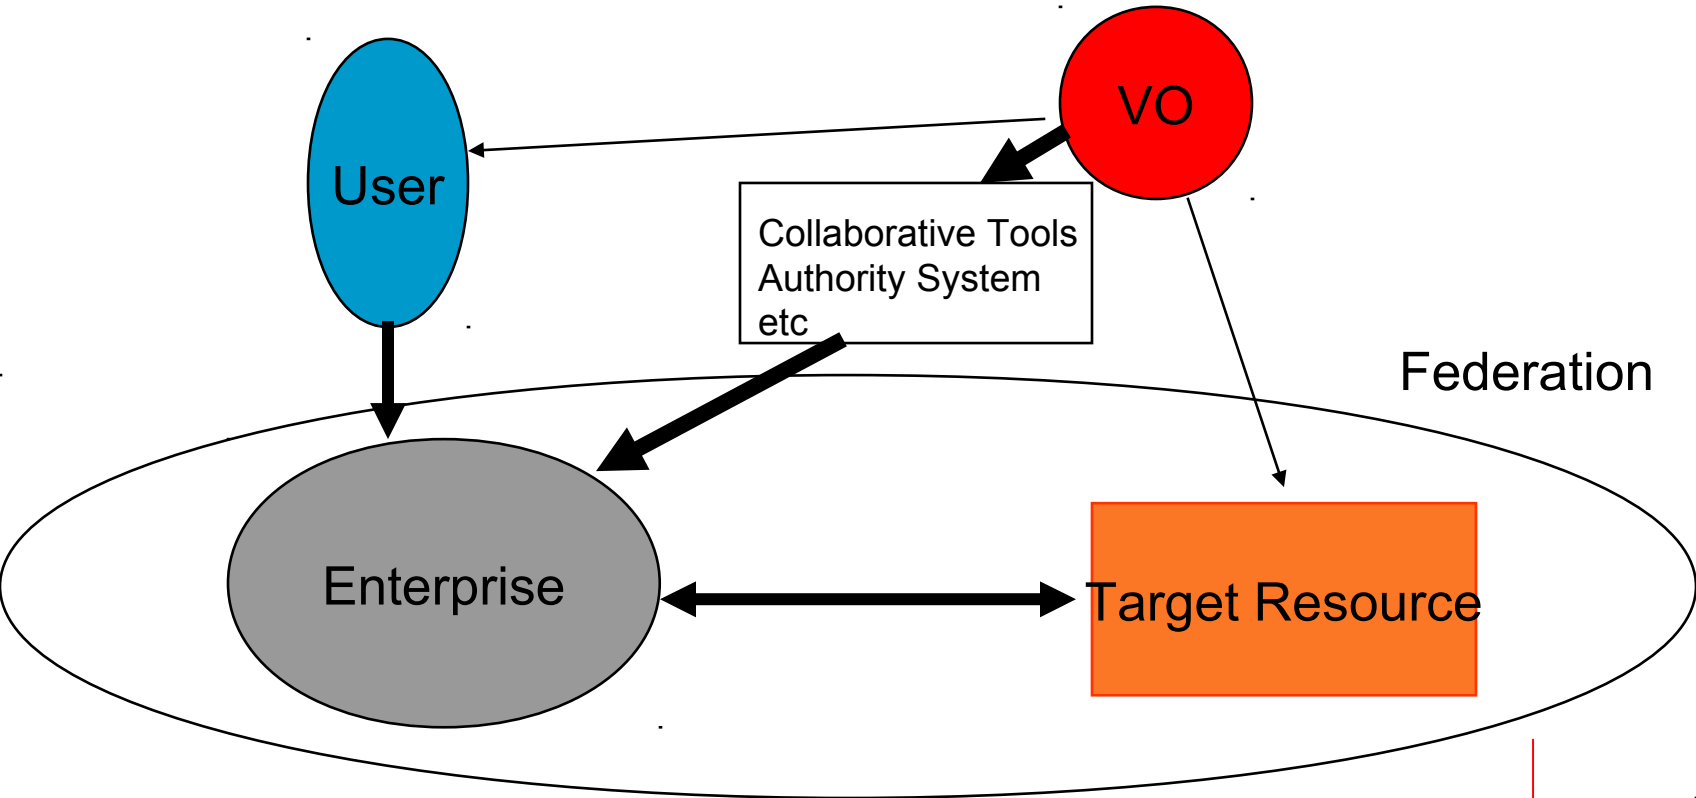
Virtual Organizations

- Geographically distributed, enterprise distributed community that shares real resources as an organization.
- Examples include team science (NEESGrid, HEP, BIRN, NEON), digital content managers (library cataloguers, curators, etc), life-long learning consortia, etc.
- On the continuum from interrealm groups (no real resource management, few defined roles) to real organizations (primary identity/authentication providers)
- Want to leverage enterprise middleware and external trust fabrics

Leveraging V.O.s Today



Leveraged V.O.s Tomorrow



Looking at the data

- Federation update
 - Federating software
 - Types of federations
 - InCommon
 - Other federations
- P2P
- V.O.'s
- Authority Systems

Shibboleth Status

- Open source, privacy preserving federating software
- Being very widely deployed in US and international universities
- Target - works with Apache(1.3 and 2.0) and IIS targets; Java origins for a variety of Unix platforms.
- V2.0 likely to include portal support, identity linking, non web services (plumbing to GSSAPI,P2P, IM, video) etc.
- Work underway on intuitive graphical interfaces for the powerful underlying Attribute Authority and resource protection
- Likely to coexist well with Liberty Alliance and may work within the WS framework from Microsoft.
- Growing development interest in several countries, providing resource manager tools, digital rights management, listprocs, etc.
- Used by several federations today – NSDL, InQueue, SWITCH and several more soon (JISC, Australia, etc.)
- <http://shibboleth.internet2.edu/>

Federations

- Associations of enterprises that come together to exchange information about their users and resources in order to enable collaborations and transactions
- Enroll and authenticate and attribute locally, act federally.
- Uses federating software (e.g. Liberty Alliance, Shibboleth, WS-*) common attributes (e.g. eduPerson), and a security and privacy set of understandings
- Enterprises (and users) retain control over what attributes are released to a resource; the resources retain control (though they may delegate) over the authorization decision.
- Several federations now in construction or deployment

InCommon federation

- Federation operations – Internet2
- Federating software – Shibboleth 1.1 and above
- Federation data schema - eduPerson200210 or later and eduOrg200210 or later
- Becomes operational mid-April, with several early entrants to help shape the policy issues.
- Precursor federation, InQueue, has been in operation for about six months and will feed into InCommon
- <http://incommon.internet2.edu>



InQueue Origins

2.12.04

- Rutgers University
- University of Wisconsin
- New York University
- Georgia State University
- University of Washington
- University of California Shibolet Pilot
- University at Buffalo
- Dartmouth College
- Michigan State University
- Georgetown
- Duke
- The Ohio State University
- UCLA
- Internet2
- Carnegie Mellon University
- National Research Council of Canada
- Columbia University
- University of Virginia
- University of California, San Diego
- Brown University
- University of Minnesota
- Penn State University
- Cal Poly Pomona
- London School of Economics
- University of North Carolina at Chapel Hill
- University of Colorado at Boulder
- UT Arlington
- UTHSC-Houston
- University of Michigan
- University of Rochester
- University of Southern California

InCommon Management

- Operational services by I2
 - Member services
 - Backroom (CA, WAYF service, etc.)
- Governance
 - Executive Committee - Carrie Regenstein - chair (Wisconsin), Jerry Campbell, (USC), Lev Gonick (CWRU), Clair Goldsmith (Texas System), Mark Luker (EDUCAUSE), Tracy Mitrano (Cornell), Susan Perry (Mellon), Mike Teetz, (OCLC), David Yakimischak (JSTOR).
 - Project manager – Renee Frost (Internet2)
- Membership open to .edu and affiliated business partners (Elsevier, OCLC, Napster, Diebold, etc...)
- Contractual and policy issues being defined now...
- Likely to take 501(c)3 status

Trust in InCommon - initial

- Members trust the federated operators to perform its activities well
 - The operator (Internet2) posts its procedures, attempts to execute them faithfully, and makes no warranties
 - Enterprises read the procedures and decide if they want to become members
- Origins and targets trust each other bilaterally in out-of-band or no-band arrangements
 - Origins trust targets dispose of attributes properly
 - Targets trust origins to provide attributes accurately
 - Risks and liabilities managed by end enterprises, in separate ways

InCommon Trust - ongoing

- Use trust ↔ Build trust cycle

- Clearly need consensus levels of I/A
- Multiple levels of I/A for different needs
 - Two factor for high-risk
 - Distinctive requirements (campus in Beijing or France, distance ed, mobility)
- Standardized data definitions unclear
- Audits unclear
- International issues

Trust pivot points in federations

- In response to real business drivers and feasible technologies

increase the strengths of

- Campus/enterprise identification, authentication practices

- Federation operations, auditing thereof

- Campus middleware infrastructure in support of Shib (including directories, attribute authorities and other Shib components) and auditing thereof

- Relying party middleware infrastructure in support of Shib

- Moving in general from self-certification to external certification

Balancing the operator's trust load

- InCommon CA
 - Identity proofing the enterprise
 - Issuing the enterprise signing keys (primary and spare)
 - Signing the metadata
- InCommon Federation
 - Aggregating the metadata
 - Supporting campuses in posting their policies

InCommon Federation Operations

- InCommon_Federation_Disaster_Recovery_Procedures_ver_0.1
 - An outline of the procedures to be used if there is a disaster with the InCommon Federation.
- Internet2_InCommon_Federation_Infrastructure_Technical_Reference_ver_0.2
 - Document describing the federation infrastructure.
- Internet2_InCommon_secure_physical_storage_ver_0.2
 - List of the physical objects and logs that will be securely stored.
- Internet2_InCommon_Technical_Operations_steps_ver_0.35
 - This document lists the steps taken from the point of submitting CSR, Metadata, and CRL to issuing a signed cert, generation of signed metadata, and publishing the CRL.
- Internet2_InCommon_Technical_Operation_Hours_ver_0.12
 - Documentation of the proposed hours of operations.

InCommon CA Ops

- CA_Disaster_Recovery_Procedure_ver_0.14
 - An outline of the procedures to be used if there is a disaster with the CA.
- cspguide
 - Manual of the CA software planning to use.
- InCommon_CA_Audit_Log_ver_0.31
 - Proposed details for logging related to the CA.
- Internet2_InCommon_CA_Disaster_Recovery_from_root_key_compromise_ver_0.2
 - An outline of the procedures to be used if there is a root key compromise with the CA.
- Internet2_InCommon_CA_PKI-Lite_CPS_ver_0.61
 - Draft of the PKI-Lite CPS.
- Internet2_InCommon_CA_PKI-Lite_CP_ver_0.21
 - Draft of the PKI-Lite CP.
- Internet2_InCommon_Certificate_Authority_for_the_InCommon_Federation_System_Technical_Reference_ver_0.41
 - Document describing the CA.

InCommon Key Signing Process

- 2. Hardware descriptions
 - a. Hardware will be laptop and spare laptop with no network capabilities, thumb drive, CDRW drive, media for necessary software
- 3. Software descriptions
 - a. OS, OpenSSL, CSP, Java tools for meta data
- 4. Log into computer
- 5. Generation of the CA Private Root key and self-signing
- 6. Generation of the Metadata signing key
- 7. Generate CSR for Internet2 origin
- 8. Signing of new metadata sites and trusts files
- 9. Backup copies of all private keys and other operational backup data are generated.
- 10. Verify CD's and MD5 checksum
- 11. Write down passphrase and put in envelopes and sign envelopes
- 12. Securely store CA hardware and contents of local safe in safe
- 13. Log that these actions occurred on the log in safe and then close and lock the safe
- 14. Put thumb drive into secure db and copy data onto secure db
- 15. Take private key password archive and other contents to Private Key Password safe deposit box and record in log that this was done.
- 16. Take operational data archive to Operation Data safe deposit box and record in log that this was done.

InCommon Process Tech Review

- As a technical review group, we, the undersigned, reviewed the processes and the following components documenting the operations of InCommon, and discussed them with the Internet2 Technical and Member Activities staff. To the best of our knowledge and experience, with no warranty implied, we believe the operational processes and procedures Internet2 provided are acceptable to begin the operations of InCommon.
 - Scott Cantor, OSU
 - Jim Jokl, UVa
 - RL Bob Morgan, UW
 - Jeff Schiller, MIT

The potential for InCommon

- The federation as a networked trust facilitator
- Needs to scale in two fundamental ways
 - Policy underpinnings need to move to normative levels among the members; “post and read” is a starting place...
 - Inter-federation issues need to be engineered; we are trying to align structurally with emerging federal recommendations
- Needs to link with PKI and with federal and international activities
- If it does scale and grow, it could become a most significant component of cyberinfrastructure...

Beyond web services...

- Federated security services
 - Collaborative incident correlation and analysis
 - Trust-mediated transparency and other security-aware capabilities
- Federated extensions to other architectures
 - Lionshare project for P2P file sharing
 - IM
 - Federated Grids

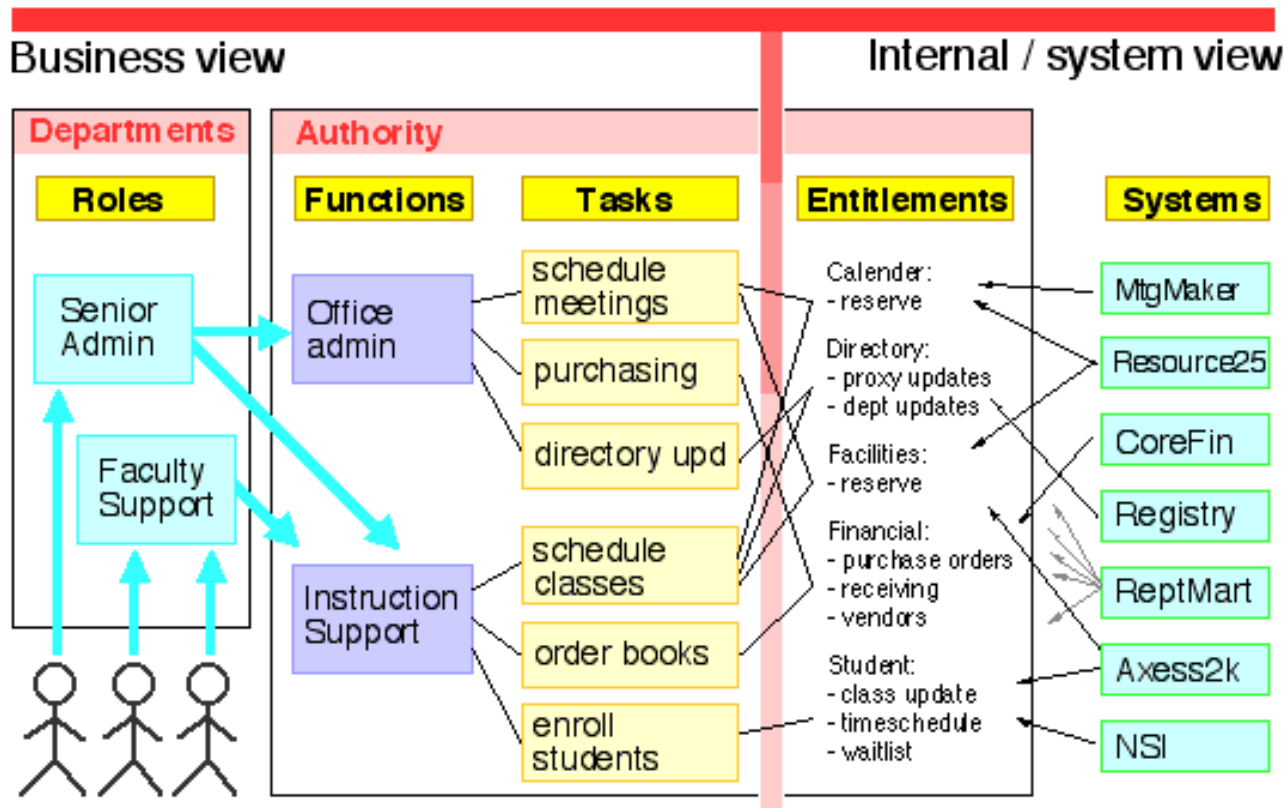
P2P arch over federated trust -Lionshare

- P2P file sharing application that is:
 - Enterprise-based – uses authentication and campus directory and resource discovery
 - Federated – works between institutions, using local authentication and authorization
 - Learning object oriented – meta-data based; linked to digital repositories, courseware, etc.
- Developed at Penn State University, now being extended with assistance from Mellon Foundation, Internet2, OKI, Edusource
- URL is <http://lionshare.its.psu.edu/main/>

Virtual organizations

- Need a model to support a wide variety of use cases
 - Native v.o. infrastructure capabilities, differences in enterprise readiness, etc.
 - Variations in collaboration modalities
 - Requirements of v.o.'s for authz, range of disciplines, etc
- JISC in the UK has lead; solicitation is on the streets (see (http://www.jisc.ac.uk/c01_04.html)); builds on NSF NMI
- Tool set likely to include seamless listproc, web sharing, shared calendaring, real-time video, privilege management system, etc.

Stanford Authz Model



Signet Deliverables

The deliverables consist of

- A recipe, with accompanying case studies, of how to take a role-based organization and develop appropriate groups, policies, attributes etc to operate an authority service
- Templates and tools for registries and group management
- a Web interface and program APIs to provide distributed management (to the departments, to external programs) of access rights and privileges, and
- delivery of authority information through the infrastructure as directory data and authority events.

The screenshot shows a web browser window titled "Authority Manager". The browser's address bar is empty, and the navigation toolbar includes buttons for Back, Forward, Stop, Refresh, Home, AutoFill, Print, and Mail. The page content is organized into several sections:

- Header:** The "AUTHORITY Manager" logo is prominently displayed. To the right are "HELP" and "END SESSION" buttons.
- User Identification:** The name "Jennifer Vine" is shown below the logo.
- Grant authority:** A section with a "GRANT New Authority" button. Below it, text reads: "If you know exactly what authority you need to grant, start here. Use Find a Person to manage an individual's authority. Use your organization views (below) to manage authority for your organization."
- View authority:** A section with the heading "By organization". Text reads: "For organizations not listed here, use Find an Organization. [Stanford University \(AA00\)](#)".
- Tools:** A section containing three tool buttons: "Find a PERSON", "Find an ORGANIZATION", and "Acting as... myself". Each button is accompanied by explanatory text: "See what a person is authorized to do. Grant, edit, or revoke their authority. View or change their designated drivers."; "See all authority scoped to any organization. (Or choose one of your organizations at left.)"; and "Select the person for whom you will act as authority-granting proxy, or return to acting as yourself."
- My authority:** A section with the heading "My authority" and two sub-items: "Jennifer Vine" and "Jennifer Vine's 'Designated Drivers'". Below these, text reads: "People who can act for you, and vice versa. Includes acting approvers and authority-granting proxies."

At the bottom of the page, there is a navigation bar with links: | Authority Home | [Help](#) | [End Session](#) |. Below this is a copyright notice: © 2001-2003 Stanford University. All Rights Reserved. The Stanford Registries logo is also present at the bottom center. A status bar at the very bottom left shows "Link: javascript:;".

Grant Authority Wizard

The screenshot shows a web browser window titled "Authority Manager" with a navigation bar containing "Back", "Forward", "Stop", "Refresh", "Home", "AutoFill", "Print", and "Mail". The main content area features the "AUTHORITY Manager" logo and a sidebar with a tree view. A modal window titled "Grant Authority: Step 6" is open, displaying the following configuration:


- 6 Set conditions for this assignment.**
- Begins on successful completion of these prerequisites:**
 - Cost Policy Training [complete]
- Expires on Lynn McRae's departure from Stanford**
- Override expiration condition** (highlighted in a red box):
 - Effective until: [] [mm/dd/yy]
 - regardless of Stanford employment status.
- Requisitions** (highlighted in a yellow box):
 - Within these limits...
 - Approval limit: **\$10000**
 - Scope: **1003321**
- 6 With these conditions...**

At the bottom of the wizard, there is a checkbox labeled "cannot extend this privilege to another person." which is checked. Below it are buttons for "<< Back", "Finish", and "Cancel".


At the bottom of the browser window, there is a footer with the text: "© 2001-2003 Stanford University. All Rights Reserved." and the Stanford Registries logo.



© Authority Manager



Back Forward Stop Refresh Home AutoFill Print Mail

 **AUTHORITY** Manager


Jennifer Vine





Authority assigned to
 **Jennifer Vine**
 User-Interface Designer,
 Technology Strategy and Support
 Operations

 HELP  END SESSION

 [Jennifer Vine's Designated Drivers](#)  [Find another PERSON](#)

[has an acting approver]

 **Financial System GL**

Dean's Office Operations (VAHO)		 Edit Limits & Conditions
Approvals	Budget control Granted on: 23-Jan-2004 By: Kristen A Murray	Limits: Scope: AABMQ, BACWK Conditions: While at Stanford Can grant
Approvals (continued)	Requisitions Granted on: 23-Jan-2004 By: Kristen A Murray	Limits: Approval limit: \$100000 Scope: 1001044, 1003321 Conditions: While at Stanford Can grant
FYIs	FYI Expense Journals Granted on: 23-Jan-2004 By: Kristen A Murray	Limits: Scope: 1001044, 1003321 Conditions: While at Stanford Can grant
Graduate School of Business (UAAA)		 Edit Limits & Conditions
Approvals	Requisitions Granted on: 27-Oct-2003 By: Wendy Jones Via: Jennifer Vine	Limits: Approval limit: \$2000000 Scope: All projects and tasks Conditions: While at Stanford
Psychology (PXIE)		 Edit Limits & Conditions
Reporting	Reporting Non-Salary View Granted on: 13-Oct-2003 By: Kristen A Murray Via: Jennifer Vine	Limits: Scope: All projects and tasks Conditions: While at Stanford Can grant
Stanford University (AA00)		 Edit Limits & Conditions
Approvals	DPA Screening Granted on: 27-Oct-2003	Limits: n/a

Link: javascript;

Next year's talk

- Lots of if's even in the premise...
- Interfederation issues
- V.O.'s over P2P trust
- Federated and progressive PKI
- Diagnostic Hell

Inter-federation Issues

- Clearly in the cards
- Some reduction of complexity can be done up front
 - Using standard technology assessment methodologies
 - Using standard policy frameworks
- Some is not problematic
 - Different objectclasses are fine
 - Different transport (Shib, WS-*, Liberty, SAML) technologies may be fine
- Much appears hard
 - Different assessment values for authn approaches
 - Different authn requirements for similar resources
 - Different privacy policies

V.O.'s over P2P trust

- P2P trust is ubiquitous (e.g. file sharing, IRC) and unraveling (viruses, abuses)
- We'll be working on V.O.'s over federated trust
- If new P2P trust tools work (controlled, integrated with apps, etc.) then V.O.'s over P2P trust represent a huge win in scaling.

Federated and progressive PKI

- Federated
 - Thin USHER – heavy id proofing of enterprises; tight operations, little policy assertion
 - Local PKI and SAML/InCommon interrealm
- Progressive PKI
 - How can we build a PKI that allows, even facilitates, growth in trust levels?
 - N-tuples of LOA's? Stochastic LOA's?

Middleware Diagnostics Problem Statement

- The number and complexity of distributed application initiatives and products has exploded within the last 5 years
- Each must create its own framework for providing diagnostic tools and performance metrics
- Distributed applications have become increasingly dependent not only on the system and network infrastructure that they are built upon, but also each other
- When what you're selling is integration and transparency, and it doesn't work...

Private Revocation Test using Oblivious Membership Evaluation Protocol

Hiroaki Kikuchi

Tokai University

Dept. of Information Media Technology,

1117 Kitakaname, Hiratsuka, Kanagawa, 259-1292, Japan

kikn@tokai.ac.jp

Abstract This paper presents a cryptographic protocol for the authenticated dictionary, namely, an untrusted directory provides a verifiable answer to a membership query for a given element. In our protocol, a user is able to retrieve whether or not a target element belongs to a database that the directory has without revealing which element he/she wishes to know against the untrusted directory. Our protocol requires linear exponentiations to the number of elements in the database, but achieves a constant size communication complexity between a user and a directory. The privacy of query is assured under the Φ -hiding assumption introduced by Cachin.

1 Introduction

1.1 The PKI Issue

Certificate revocation is a current topic of interest in public-key infrastructure (PKI). Traditionally, a list of revoked certificates (CRL) has been used to represent the periodic distribution of revoked information. To improve the bandwidth consumption of the entire CRL transmission, some mail agents have begun supporting an online protocol for providing users the status of a target certificate alone, instead of the full CRL. The Online Certificate Status Protocol (OCSP)[6] is the standard protocol now in common use. There have been several attempts to improve the efficiency and security of CRLs. Kocher proposed a hash-tree based revocation protocol known as CRT[8], Micali presented a linear linking scheme with $O(1)$ communication cost (CRS)[7], and Naor and Nissim formalized the problem as an authenticated dictionary [9] in which a B-tree is used to balance the tree while the tree itself is skewed while updating the database.

1.2 Privacy Issues

As these online protocols are now in widespread use, a new privacy issue has arisen. The OCSP method uses the following steps. Each time a digitally signed mail is received, then the mail agent picks up a certificate from the mail and automatically sends a query to check if the certificate is revoked to a server specified in the certificate. Hence, the server, known as the *CRL distribution point*, acquires the significant statistics of the PKI – who sends a message to whom, how often, and, even worse, a digital signature, which is often used when we send significant messages whose privacy we wish to preserve the most.

1.3 Privacy Information Retrieval

To overcome the privacy issues of revoked certificates, the private information retrieval (PIR) method is a suitable technique for a user to be able to retrieve a target data item from a database while hiding the identity of the target item from the server. The notion of a PIR was introduced by Chor, Goldreich, Kushievitz, and Sudan [4], and has already improved retrieval in terms of its communication and computation costs. One of the recent results by Beimel, Ishai, and Malkin [5] archives, for a given constant, $k \geq 2$, and the number of items in a database n , a k -server protocol with $O(n^{1/(2k-1)})$ communication, and $O(n/\log^{2k-2} n)$ computations at the server. The servers, however, are considered as untrustworthy parties in the PKI model because servers must be online and, thus, have greater chance of being compromised by an intruder. Therefore, the behavior on the server side is not guaranteed to be correct. In addition, the average user may have poor computational power and narrow bandwidth, with even just one server. Thus, a single server protocol making the cost at user side as small as possible is preferable for solving the CRL distribution problem.

1.4 Our Contribution

In this paper, we present a simple solution to the problem. Given an element, $x \in X$, a user performs a membership test if x is in a subset $L = \{x_1, \dots, x_n\} \subset X$, requesting a query for a single non-trusted server who manages L steps without revealing x to the server. Our proposed protocol achieves a single server PIR with an optimal communication cost of $O(1)$ between a server and a user, and an optimum computation cost of $O(1)$ at the user side. To prevent the server from answering an improper response, a verification protocol that authorizes the answer from an authority is also provided.

2 Preliminaries

2.1 The PKI Model and Requirements

We have three types of parties: The *source* S or *Certification Authority* (CA), which is a trusted party that certifies the list of revoked certificates, *directories* D is non-trusted party who maintains the list and answers the questions that a target certificate is still active, instead of CA, and *users*, U , who wish to keep in touch with the current status of the certificates via the non-trusted D .

The CA is a source of information of revoked certificates and has the replication of the information distributed among directories. The directories of D s work as carriers of the revoked information and are thus not responsible for the integrity of the database provided from the source. In terms of security, the directories have no secret information inside so that, even if one of directories is compromised, any rebuild of PKI is not necessary. The directories have a powerful computational power, e.g., the state-of-the-art CPUs, a secure coprocessor, and broad-bandwidth connections to each party. Since the directories are widely distributed over the network, we assume the risk that some of directories might perform an analysis of the access log from the end users using the data mining techniques. A user U communicates with one of the directories and checks if a target certificate is revoked or not and examines the integrity of responses from the directory server. We assume that some of the users may have limited computational power and a poor link of limited bandwidth. (In particular, this can happen when the user is mobile and with a PDA).

Oblivious Membership Evaluation:

Let X be the universal set of identities of certificate (64-bit serial numbers are often used in actual services), and $L = \{x_1, \dots, x_n\}$ be a subset of X . The S gets L distributed among directories D . Given an element $x \in X$, U performs a membership query to D whether or not, $x \in L$ without revealing x to D .

The requirements of oblivious membership evaluation should satisfy are as follows:

1. **Privacy of query.** From a membership query of $x \in L$, D learns no information about x .
2. **Authenticity of source.** From the response from D , U verifies that the result of membership query is authorized by S and that D follows the steps properly.
3. **Efficiency.** The sizes of both query and answer should be independent of the number of PKI users, to which the size of CRL n seems to be proportional, and we want the sizes to be as small as possible. The computational costs at users should be also minimized.

2.2 Dynamic Accumulator

The RSA accumulator is proposed by Benaloh and de Mare[10], where a set of values are accumulated into a single object for which a witness that a given

value was incorporated into it is provided. Camenish and Lysyanskaya improves the RSA accumulator so that dynamic operations of insertion and deletion are feasible with independent cost from the number of values[12]. Goodrich, Tamassia, and Hasic show the pre-computations of witness reduces the computation overhead at the directory with the cost of communication consumption[11].

Informally, the RSA accumulator works as follows. The source picks strong primes p and q and publishes $N = pq$. Let L be a set of primes $\{x_1, \dots, x_n\}$, representing identities (of the revoked certificates in PKI). The source then computes *accumulator*

$$A = a^{x_1 x_2 \dots x_n} \pmod{N},$$

where a is a public constant that is relatively prime to N and publishes A together with digital signature $\sigma_S(A)$ on A . To prove an element $x_i \in L$, the directory computes *witness*

$$A_i = a^{x_1 \dots x_{i-1} x_{i+1} \dots x_n} \pmod{N}.$$

The user verifies witness by $A_i^{x_i} \pmod{N} \stackrel{?}{=} A$. Under the strong RSA assumption[12], the directory, which does not have the knowledge of factorization of N , is able to compute the witness A_i only when x_i belongs to L .

2.3 Φ -Hiding Assumption

Cachin present an efficient secure auction protocol that an oblivious party blindly compares two inputs bit-by-bit under the the ϕ -hiding assumption (Φ HA)[13]. Informally, the Φ HA states that it is computationally infeasible to decide whether a given prime divides $\phi(N)$, where m is a composite number of unknown factorization.

We say modulus m *hides* a prime p if N is a composite number $p'q'$ such that $p' = 2pp_1 + 1$ and $q' = 2q_1 + 1$ with primes p_1, q_1 . Note that N hides p if and only if $p|\phi(N)$. The Φ HA states that, for a randomly chosen $N \in Z_N^*$ and primes p_0, p_1 such that N hides p_0 but does not hide p_1 , the (N, p_0) and (N, p_1) is computationally indistinguishable.

An integer x is a p -th *residue* modulo m if there exists an α such that $\alpha^p = x \pmod{m}$. Let $R_N(p)$ denote a set of all p -th residues in Z_N^* . Then, note that only the party that knows the factorization of N and thus $\phi(N)$ is able to test if any given integer is a p -th residue by

$$a^{\phi(N)/p} \equiv 1 \pmod{m},$$

which holds if a is a p -th residue modulo m .

2.4 Proof of Conjunctive Knowledge

Cramer, Cramer, Damgard, and Schoenmakers presents an efficient zero-knowledge proof of conjunctive propositions [1]. By $PK\{(\alpha) : y_1 = g_1^\alpha \wedge y_2 = g_2^\alpha\}$, we denote a proof of knowledge of discrete logarithms of elements y_1 and y_2 to the

bases g_1 and g_2 . Selecting random numbers r_1 and $r_2 \in Z_q$, a prover sends $t_1 = g_1^{r_1}$ and $t_2 = g_2^{r_2}$ to a verifier, who then sends back a random challenge $c \in \{0, 1\}^k$. The prover shows $s = r - cx \pmod{q}$, which should satisfy both $g_1^s y_1^c = t_1$ and $g_2^s y_2^c = t_2$.

3 Oblivious Membership Evaluation

3.1 Overview

Our construction is based on Φ HA in order for users to blindly query a membership to a directory that has the list L . A user generates a modulus m that hides a prime x specifying the identity of a given certificate, and then sends a query consisting of non x -th residue c . The directory D then raises c to the power of all primes in S modulo m and sends the answer back to U , who then performs an x -th residue test using secret knowledge of factorization of m . In addition, we need a verification protocol to prevent a dishonest directory from cheating users. The witness in RSA accumulator cannot be applied here because the directory does not know which element is to be tested. Instead, we employ a zero-knowledge proof technique to show that the directory has raised a base to the power exactly the same exponents to that used by accumulator A .

3.2 Accumulator Setup

We begin with a set up protocol in which a source S notifies to the directories the list of currently revoked certificates.

1. The S picks strong primes P and Q and publishes $N = PQ$. For the list of revoked certificates $L = \{x_1, x_2, \dots, x_n\}$, where x_i are small primes corresponding identities of revoked certificates, S computes accumulation

$$A = a^{x_1 x_2 \dots x_n} \pmod{N},$$

where a is a public constant that is relatively prime to N and publishes L, A, a together with a digital signature $\sigma_S(A, a, t)$, where t is the current time interval.

2. On receiving the list L and accumulator A periodically, every directory D updates the current (at a time t) database by L after it verifies the digital signature and accumulator $a^{x_1 x_2 \dots x_n} = A \pmod{N}$.

3.3 Membership Test

Given a certificate to be examined, a user performs the following membership test protocol with one of the directories.

1. Given a target certificate specified by prime x , U chooses strong primes p and q such that $m = pq$ hides prime x . U picks an integer c that is not

p -th residue modulo m . U sends a query of the form (c, m) to one of the directories, D .

- Then, D computes an answer

$$z = c^{x_1 x_2 \cdots x_n} \pmod{m}$$

and responds to U the answer z together with the accumulator A and digital signature $\sigma_S(A, a, t)$.

- Finally, U locally performs the membership test

$$z^{\phi(m)/x} \pmod{m} = \begin{cases} 1 & \text{if } x \in L, \\ 1^{1/x} & \text{otherwise.} \end{cases}$$

Note that answer z becomes the x -th residue when there is an element in L that is equal to the target x .

3.4 Authenticity of the Source

To prevent a dishonest D from cheating users with improperly computed z , we require D to provide the proof of accumulating every element L into A by the form

$$PK\{\beta : a^\beta = A \wedge c^\beta = z\},$$

where private information β is ℓ defined by $\ell = x_1 x_2 \cdots x_n$ (note that this is not a modular multiplication), for which both $z = c^\ell$ and $A = a^\ell$ are satisfied. In other words, ℓ is a witness for which accumulator A is consistent with the answer z . Since D does not know the factorization of N nor m , we need the modified version of the proof of conjunctive knowledge mentioned in Section 2.4.

- The D randomly picks r that is properly large (but is less than N and m) and computes

$$T = a^r \pmod{N}, \quad V = c^r \pmod{m}.$$

For T and V , D applies a secure cryptographic hash function H with properly large range to obtain a challenge $d = H(T||V)$, and computes (not modular arithmetic)

$$s = r + d\ell$$

and sends the proof (T, V, s) to U .

- Then, U computes $d = H(T||V)$ and verifies that

$$\begin{aligned} a^s / A^d &\stackrel{?}{=} T \pmod{N}, \\ c^s / z^d &\stackrel{?}{=} V \pmod{m}. \end{aligned}$$

4 Evaluation

4.1 Security

Under the assumption of a secure digital signature scheme used by the source, the accumulator A at the time t is unable to be forged. Consider a dishonest directory that is trying to manipulate z to z' so that the membership test will fail for z' when x is in L . To convince users that the answer was correctly computed, the directory has to predict s that satisfies the above-mentioned equations for proof of knowledge. The probability of passing the test is negligibly small.

4.2 Privacy

If a malicious directory is able to determine which prime is hidden by a given m and c , it can immediately distinguish two composite numbers m_0 and m_1 that hide distinct primes, which contradicts the Φ -hiding assumption. Therefore, D is not able to learn the target x under the Φ HA. Moreover, D does not even know the result of the membership test at all.

4.3 Efficiency

The proposed scheme has the following performance:

- a size of query (c) sent from user to directory is $|m|$;
- a size of answer (z) sent from directory to user is $|m| + |N| + |\sigma|$ (without proof of knowledge);
- a size of proof (T, V, s) is $O(n)$ (since the magnitude of ℓ is linear to n);
- a number of modular exponentiations at the user is 1;
- a number of modular exponentiations at the directory is n .

Without the knowledge of $\phi(m)$, the size of ℓ increases with the number of elements in L ; thus, the verification at the last step in the scheme requires $O(n|m|)$ modular multiplications, which is impractically heavy when n is too large.

One more inefficiency we should address is the key generation cost to the user, who should always pick a new modulus that hides the given prime.

5 Conclusions

We have presented a protocol for oblivious membership evaluation using the Φ -hiding assumption. The proposed protocol is efficient in terms of directory-and-user communication with $O(1)$, preserves the privacy of a query as to which certificate is to be examined, and provides verification steps that result in the membership query being correctly computed. Future studies include an efficient

verification independent of n and an improvement of n -size modular exponentiations at the directory.

References

- [1] R. Cramer, I. Damgard, and B. Schoenmakers, "Proofs of partial knowledge and simplified design of witness hiding protocols," in Proc. of *CRYPTO '94*, pp.174-187, 1994.
- [2] J. Camenisch, and M. Michels, "Proving in zero-knowledge that a number is the product of two safe primes," in Proc. of *EUROCRYPT '99*, pp. 107-122, 1999.
- [3] R. Cramer, R. Gennaro, and B. Schoenmakers, "A secure and optimally efficient multi-authority election scheme," *EUROCRYPT 1997*.
- [4] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan, "Private information retrieval," in Proc. 36th IEEE Symposium on Foundations of Computer Science (FOCS), 1995.
- [5] Amos Beimel, Yuval Ishai, and Tal Malkin, "Reducing the servers computation in private information retrieval: pir with preprocessing," in Proc. *CRYPTO '00*, LNCS vol. 1880, pp. 550, 2000.
- [6] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams, "X.509 Internet public key infrastructure online certificate status protocol - OCSP," Internet RFC 2560, 1999.
- [7] S. Micali, "Efficient certificate revocation", Technical Report TM-542b, MIT Laboratory for Computer Science, 1996.
- [8] P. Kocher, "On certificate revocation and validation," in Proc. of Financial Cryptography'98, *Springer LNCS 1465*, pp. 172-177, 1998.
- [9] M. Naor, and K. Nissim, "Certificate revocation and certificate update," in Proc. of Seventh USENIX Security Symposium, pp. 217-228, 1998.
- [10] J. Benaloh, and M. de Mare, "One-way accumulators: A decentralized alternative to digital signatures," in Proc. of *EUROCRYPT*, LNCS vol. 839, Springer, pp. 216-233, 1994.
- [11] M. Goodrich, R. Tamassia, and J. Hasic, "An efficient dynamic and distributed cryptographic accumulator," in Proc. of Information Security Conference (ISC 2002), LNCS Vol. 2433, Springer, pp. 372-388, 2002.
- [12] J. Camenisch, and A. Lysyanskaya, "Dynamic accumulators and application to efficient revocation of anonymous credentials," in Proc. *CRYPTO 2002*, 2002.

- [13] C. Cachin, "Efficient private bidding and auctions with an oblivious third party," ACM Conference on Computer and Communications Security (CCS), pp. 120-127, 1999.

“Dynamic Bridge” Concept Paper

Ken Stillson
Mitretek Systems

Written for presentation to the
3rd Annual PKI R&D Workshop

Table of Contents

Problem Background	3
Solution Background	5
Dynamic Bridge Concept.....	6
Proposed Benefits	9
Challenges and Unresolved Issues.....	10
Request for Feedback.....	11
References.....	12

Abstract

Current cross-certificates based PKI trust mechanisms suffer from a scaling problem. Even given the topological simplification of bridge CAs, as cross certificate meshes grow in size and complexity, the number of possible routes between points increases very quickly, and the time required for path discovery can increase beyond a tolerable delay for real-time operation. This paper proposes a “dynamic bridge,” which is an automatically created transformation of a cross certificate topology, designed to reflect the same trust arrangements and constraints, but in a simplified structure. Creation of dynamic bridges should not require centralized coordination or infrastructure, and use of them for speed-enhanced validation should require clients to implement only a subset of standard path discovery and validation logic.

Problem Background

The standards that govern PKI, primarily [X509] and [RFC3280], envision a mechanism for a recipient, or “relying party” in one PKI domain to accept credentials from a sender or signer in another. The recipient has one or more “trust anchors.” These are certificate authorities (“CAs”) that the recipient trusts completely. These CAs can create cross-certificates, which indicate other CAs that this CA trusts. This process can be repeated multiple times, resulting in a chain of trust from the trust anchor to the sender’s certificate. [pathbuild]

This process involves three distinct areas: “path discovery,” “object location,” and “path validation.”

Path discovery entails finding the possible chain(s) of certificates between the sender and the trust anchor(s). Path discovery would be challenging enough even if all the certificates in the world were immediately available to select from. However, generally the only inputs to path discovery are the end-points: the sender’s certificate, available because it is included with the signed message being validated, and the trust-anchor(s), which are part of the relying party’s configuration. Path discovery therefore involves an iterative process sniffing out each “next possible link” – building the chain one link at a time.

Object location is the challenge of retrieving the certificates and cross certificates needed to feed the path discovery algorithm. Object location suffers from competition between several different mechanisms, none of which are very mature, and each of which assume they are the global solution. The separate mechanisms do not easily build upon each other. This challenge is not fatal, as the software of a relying party can support all of the contending mechanisms.

Path validation takes a candidate chain created by path discovery, and confirms that all the rules of trust transfer are followed within the chain. For example, cross-certificates can stipulate constraints that add requirements to the overall chain, or to parts of it some distance away from the certificate adding the constraint. Path validation checks all the constraints of each certificate against the others. Path validation also generally includes an “on-line status check” to confirm that no certificate in the chain has been revoked. Path validation is a computationally expensive process, but is well defined and reasonably well understood.

The focus of this paper is a scaling problem with path discovery.

Although developed independently, the concept is an extension of [Sun1], which envisions hierarchical root CAs issuing certificates directly to their n-level subordinates, thus flattening the hierarchical structure. The dynamic bridge extends this concept into the space of multi-domain PKIs linked by cross certificates, and handles the complexity of policy and constraint mapping introduced by such an extension.

In figure 1, a path discovery algorithm starts with the sender's certificate¹. An object location mechanism should easily find CA 1, as CA 1's name is stated in the sender's certificate. The next step will be for the discovery algorithm to ask the location mechanism to find all certificates issued to CA 1. This could result in any number of certificates, indicated by the multiple green arrowheads around CA 1. As we have the picture already laid-out, we can see that the link to CA 2 is the correct choice. However, there is no way for the discovery algorithm to know this. It may well select the link that leads into cloud X, and one can imagine that if cloud X contains a large and complex mesh, it may be some time before the path discovery algorithm realizes it made a wrong turn at CA 1, and tries the alternative path leading to CA 2.




Figure 1

One response to this problem has been the implementation of PKI “bridge CAs.” [pathbuild][FBCA]. A bridge CA is like a trust “hub”; it re-organizes the cross-certificate “mesh” topology into a star shape, which shortens the number of links in chains.




Figure 2

¹ There is an alternate technique [RFC3280] that begins at the trust anchor(s) and builds the path towards the sender. This difference is not relevant to the scaling problem that the example is building towards.

If an all-encompassing central bridge linked the entire world, there would be no path discovery problem. The cross-certificate from each trust anchor to the bridge would contain enough information to be located immediately from the bridge.

However, it has become clear that there will be no global bridge in the near future. No organization appears to have the combination of desire, funding, expertise, and ubiquitous acceptance that would be required. Rather, individual arenas are creating bridges that cover their natural scope. For example, the US Federal government has established the Federal Bridge CA [FBCA], Canada has a national bridge underway, and EDUCAUSE (a consortium of educational institutions) has created the Higher Education Bridge CA [HEBCA], etc.

These bridges are slowly being linked together. The result will likely be a cluster of large bridges, surrounded by a constellation of smaller bridges, surrounded by individual PKI domains. While this arrangement is an improvement on an unstructured mesh, in that the *average* path length will be lower, the problem discussed above in figure 1 remains. In fact, in figure 1, if CA1 and CA2 were both bridges, the number of possible “wrong” routes would likely increase drastically, compared to individual CAs.

Basically, path discovery suffers from a lack of a “sense of direction.” [PKI concepts]

Solution Background

A common approach to resolve the path discovery sense of direction problem involves a process scanning the entire cross-certificate mesh, and pre-processing the results in some way to make discovery of a specific path faster once the actual endpoints are known.

However, there is a general complication to this technique. One cannot “cache” pre-validated *partial* paths. This is because of the ability of any certificate in a chain to add constraints applying to other (non-neighboring) certificates. Specifically, if a chain from CA1 ↔ CA2 ↔ CA3 is found, and is believed to be a common component to complete chains, it would make sense to cache it as an available component for building larger chains. However, once CA4 is added, CA4 may contain a property forbidden by a constraint in CA2, or vice-versa. As this is true all the way through to the end-user certificate, no partial path is safe from elimination by constraints from certificates not included in the partial path.

This is a complication rather than a fatal flaw as it can be resolved by careful separation of path discovery and path validation. In other words, partial paths can be cached as “discovered,” so long as path validation is performed on the complete path once assembled, to make sure that constraints of the additional certificates are respected.

An example of a path discovery assistance system is the “intermediate store solution,” envisioned and implemented in proof-of-concept by the FBCA’s “path discovery and validation working group (PD-VAL).”²


This intermediate store solution uses a “cross certificate spider” [HHSspider] to map out entire cross certificate topologies, and retrieve all the directly linked certificates and cross-certificates. This program essentially takes on the object location challenge, and retrieves all objects possibly needed to a server. The certificates are then stored into a distributable form (a PKCS7 file), and published to a web-server. A program running on end-user Microsoft workstations then regularly downloads the PKCS7 file, and adds its contents to the Microsoft Windows “intermediate store.” This removes the iterative piece-by-piece part of path discovery, leaving only the problem of selecting a valid chain given the pre-collected universe of certificates, a capability that is built-in to most modern versions of Windows.³

This approach shows some promise, in homogenizing differences between versions of Windows (its original purpose), solving the object location problem, and simplifying the path discovery problem. [CAI-POC] However, like most pre-caching solutions, it requires proprietary software running on the desktop to utilize the cache – in this case, the program that downloads and installs the PKCS7 cache file.

Dynamic Bridge Concept

The idea of the dynamic bridge is to actively search out paths with a path length greater than 1 hop, and “condense” them by creating direct cross-certificates to reduce the path-length to one.

For example we consider the path




Here, X, Y, and Z are constraints, policies, policy mappings, and other attributes that effect a path’s validity.

² “PD-VAL” was established by the FBCA Operational Authority to research challenges on bridge-aware path discovery. PD-VAL members include representatives from NIST, the Federal PKI technical working group, various Federal agencies, PKI vendors, and the author.

³ Some versions of Windows do have the capability to perform iterative path discovery, but they rely on an object location solution that makes assumptions about certificate “AIA fields” that are not universally followed. By separating the object location function into the “cross-certificate spider,” which supports multiple object location mechanisms, this removes the Windows AIA requirement.

Dynamic bridge infrastructure would transform the above situation as follows:



Specifically– A new CA, “Dyn(A),” with its own self-signed certificate, has been created, and that new CA has issued a series of unidirectional cross certificates. The new cross certificates reflect exactly the same trust arrangements between their subjects and Dyn(A) as the initial CA, but they have been “flattened” to combine intermediate hops. All trust is transferred directly from Dyn(A) to the targets with new 1-hop cross certificates.

This “condensation” of multiple hops into a single one is similar to [Sun1], except that rather than flattening hierarchical chains, the condensed hops here include cross-certificates between distinct PKI domains. This creates a significant new complication – trying to condense the constraints of the cross certificates.

In the case of the trust transfer from A -> CA2, with the constraint X, nothing has changed. However, in the case of the path A -> CA2 -> CA1 with the constraint X between A and CA2, and the constraint Y between CA2 and CA1, this path of length two has been reduced to a path of length one, directly from Dyn(A) to CA1, with the new constraint “ $Y \oplus X$ ”, which is the ordered combination of constraints Y and X.

For example, if Y requires policy1, and X requires policy2, then $Y \oplus X$ would require both policies in a single constraint. If Y maps policy1 to policy2, and X maps policy2 to policy3, then $Y \oplus X$ map policy1 directly to policy3. It is asserted that all the constraints possible in cross certificates can be combined by a fully implemented \oplus function, and that therefore the types of certificates and chains that can be condensed is not limited.

A dynamic bridge infrastructure would automatically create Dyn(A) by using a cross certificate spider to locate all paths that lead back to A. Once this process is completed, a user that previously used A as their trust anchor would change their trust anchor to Dyn(A). They would now find that all paths leading back to A are now immediately available directly from Dyn(A) with path length 1.

Furthermore, it is proposed that all cross certificates created with the \oplus function also add a path length constraint indicating that at most one further hop is permitted. This would effectively prevent iterative path discovery. If a user uses only Dyn(A) as a trust anchor, then the process of path discovery is reduced to selecting the cross certificate from the issuer of the sender's certificate to the trust anchor.

Walking through a typical path discovery algorithm: in typical implementations, discovery starts with the sender's certificate, which contains the name of its issuer. An object location query is executed for all certificates whose subject matches the issuer of the sender's certificate. In this case, only CA1 will be returned. CA1 is not a trust anchor, so the algorithm iterates. An object location query is run for all certificates whose subject is CA1. This will return numerous certificates, including the one issued directly from the dynamic bridge. As the dynamic bridge is a trust anchor, path discovery is concluded. The key is that the path discovery algorithm never had to make a "guess" as to which certificate to select. The very first query that returned multiple certificates including a direct link to the trust anchor. This avoids the need for a "sense of direction."⁴

There is an interesting implication for revocation checking. The dynamic bridge path has "short circuited" CA2. A correctly implemented \oplus function would ensure that CA2's constraints are respected, however in the original topology, CA2 also has the capability to revoke its cross certificate to CA1, thus breaking the path. The dynamic bridge path does not pass through CA2, and thus removes CA's revocation capability.

There is a solution. The expiration date for the dynamic bridge's certificate from Dyn(A) to CA1 should be set to the earliest CRL "next update" time for any of the CA's that have been "flattened" from the path. i.e. The earliest CRL update time has become the cross certificate expiration time. In this way, a dynamic bridge cross certificate is valid only up until the time that a CRL update could have invalidated part of the consolidated path.

The dynamic bridge must continuously re-generate its cross certificates as they expire, and obviously will re-perform full path validation before re-consolidating paths, thus giving the intermediate CA's the opportunity to revoke paths.

Clearly the dynamic bridge is going to be a busy system. Its internal database conceivably consists of certificates for every CA in the world, it must continuously scan for new CA's (new paths to add), and perform the above re-validation of existing paths each time a CRL expires. However, a clever implementation could filter on changes to previous queries to detect additions, and queue re-validation carefully for only expiring paths. The assertion is made that this implementation is feasible.

⁴ Note that it is possible that multiple paths existed through the original mesh between the sender's CA and the trust anchor. In this case, there would be multiple single-hop cross certificates from dynamic bridge to the issuer's CA. Although there is a "choice" as to which of these multiple paths will be selected, this does not change the claimed advantage. Whichever single hop cross certificate is selected by the path processor, it results in immediate path to the trust anchor with no further iteration, there is no "heading off in the wrong direction" regardless of which certificate is selected.

Proposed Benefits

It is believed that this arrangement would lead to benefits in each of the three primary areas of cross certificate-based trust building.

For path discovery, the need for iterative discovery is removed. The path length constraints added by the \oplus function ensure that no “wild goose chases” will occur by the path discovery algorithm taking a wrong turn; as no “turns” are allowed.

For path validation, the process of validation has become simpler, both because path lengths are reduced, and because the cumulative effects of the chained constraints have been pre-calculated.

For object location, the dynamic bridge process has already collected and consolidated all the objects that the relying party will require. Assuming the dynamic bridge’s cross-certificates are all stored in a single directory, the recipient’s object location system needs only to be directed to that location.⁵ If the dynamic bridge’s object location mechanism supports multiple of the competing retrieval techniques (e.g. DN searching against X.500, AIA, LDAP referral trees, etc), the dynamic bridge’s user’s software need only support the mechanism that leads it to the dynamic bridge’s consolidated repository.

Construction of a dynamic bridge requires no particular privilege, nor the cooperation of the mesh participants, other than CAs posting their cross-certificates into locations that the dynamic bridge’s object location techniques can find. Any enterprise, or even end-users, can establish their own dynamic bridge(s). If multiple trust anchors are in use, either multiple independent dynamic bridges can be constructed (if selection as trust anchors must be separately selectable), or a dynamic bridge could merge multiple meshes by seeding the “condensing” process from multiple trust anchors.

While only those that utilize the dynamic bridge’s trust anchor will receive the above benefits, the existence of the dynamic bridge is non-harmful to those that do not utilize it.

Finally, utilization of a dynamic bridge does not require any specialized software. Any standards-compliant path discovery system will be able to gain the advantages of a dynamic bridge just by re-selecting their trust anchor(s). In-fact, only a small subset of the standards-required capabilities are needed; some PKI software that is not fully compliant now (due to lack of iterative path discovery, or limited object location capabilities) would be made fully capable by using a dynamic bridge.

⁵ Another interesting possibility is that the dynamic bridge could set AIA and/or SIA fields in the cross certificates that it generates. An example of a possible advantage -- the dynamic bridge will have multiple object location techniques that it may utilize when searching the cross certificate mesh, storing the technique that worked in an SIA field of the cross certificate could simplify object location of the target certificate for path building techniques that start with the trust anchor and work towards the sender.

Challenges and Unresolved Issues

Firstly, while the dynamic bridge does not create new key pairs⁶, it does automatically create cross certificates which its users will trust. Frequently, CAs are kept “offline” to improve security, but due to the automated and continuous nature of its processing, the dynamic bridge must be “online.”⁷

Secondly, the entire concept hinges on the \oplus function – the ability to take a series of constraints in separate certificates along a path and condense them into a single set of constraints, stored in a single X.509 compliant cross certificate constraint. Intuitively, this should be possible. A sample “permitted and forbidden sub-trees” algorithm is specified in [RFC3280], and although that algorithm is designed to be run iteratively during path construction, it should also be possible to run it during the cross certificate crawl. Furthermore, it appears that the constraints specification system is adequately expressive that transitive combinations of constraints can be simplified to a single constraint, but until \oplus is successfully implemented, caution is needed.

Thirdly, Microsoft Window’s build-in path discovery system (“CAPI”) remains a challenge with respect to object location. CAPI does not support specification of a “default directory” or an “AIA⁸ of last resort,” which could be used to point to the dynamic bridge’s directory. CAPI builds from the sender towards the trust anchor, so addition of AIA fields to dynamic bridge certificates does not help, as the sender’s certificate’s AIA will not lead to the relying party’s dynamic bridge directory. Again, this could be overcome by loading the dynamic bridge output into the relying party’s intermediate store, but this would require proprietary software on the desktop.

Finally, there is an issue with respect to the object location algorithm used by the spider process that builds the dynamic bridge. The crawl must start at the known trust anchor(s), and spread outwards. This direction is “the hard way” for object location.

When AIA fields are not present, the object location problem is generally solved via an LDAP search for DN’s matching the subject field of the desired objects. When building a path from the sender’s certificate towards the trust anchor(s), each certificate contains the DN of it’s issuer, so the next possible steps can be queried by searching for certificates with the subject that matches the issuer of the current DN. However, certificates do not have an “issuee” field, so this technique cannot proceed in the opposite

⁶ The creation of cross certificates involves signing an existing key with an existing key, it does not generate new key pairs.

⁷ Technically, only a border directory containing the cross certificates must be fully on-line. The dynamic bridge must be able to push cross certificates onto its border directory, but other than this action, can be well isolated. While an “air gap” around CA’s is “better,” it is not unusual for production CA’s to have a live one-way connection to their border directories.

⁸ AIA stands for “authority information access,” and in this context, gives a location to obtain all certificates whose subject matches the issuer of that certificate. AIA fields may be used by object location algorithms to obtain the “next step” in path discovery. An “AIA of last resort” specifies a general technique for finding any certificate’s issuers given it’s DN. Use of this type of mechanism is one of the competing object location solutions referred to in the first section.

direction. [X509] defines the “SIA” extension (the converse of the AIA field) for this purpose, but SIA is not widely populated.

There is a saving grace – bi-directional trust. When A issues a cross certificate to B, the matching reverse cross certificate is usually issued by B to A. The first iteration of the “subject that matches the issuer” technique will locate the reverse certificate, thus revealing the DN of the CA one step further from the trust anchor. The next iteration of the “subject that matches the issuer” technique will then return the forward cross certificate, allowing the crawl to spread outwards. This procedure has been shown to work in an implemented spider [HHSpider], but it does rely on bi-directional trust (or SIA fields) to discover the entire mesh.

Request for Feedback

Mitretek Systems is presenting the dynamic bridge concept to the PKI community without intent to assert patent protection, in hopes that its utility may be assessed, and discussions started concerning the possibility of implementation.

Those interested in providing feedback, or joining discussions on the topic, are asked to contact the paper’s author, Mr. Ken Stillson of Mitretek Systems, at stillson@mitretek.org, or 703-610-2965. If there is sufficient interest, a mailing list or similar discussion mechanism will be established.

References

- [CAI-POC] K. Stillson, *HHS CAI Proof-of-Concept Results*. Prepared by Mitretek Systems for the Dept. of Health and Human Services, Sept. 2003
[Available upon request from HHS]
- [FBCA] CSRC/NIST *Federal Bridge Certification Authority*
See <http://csrc.nist.gov/pki/fbca/welcome.html>
- [HEBCA] Higher Education Bridge Certification Authority, web site.
see <http://www.educause.edu/hebca/>
- [HHSspider] K. Stillson *The Certificate Spider – A Simplified Technical Description*
Mitretek Systems. June, 2003. Prepared for Mark Silverman of the
department of Health and Human Services
[Available upon request from HHS]
- [Pathbuild] M. Cooper, Y. Dzambasow, P. Hesse, S. Joseph, R. Nicholas,
Internet X.509 Public Key Infrastructure: Certification Path Building
(Internet Draft) IETF PKIX Working Group. December 2003. See
<http://www.ietf.org/internet-drafts/draft-ietf-pkix-certpathbuild-03.txt>
- [PKI concepts] K.D. Stillson *Public Key Infrastructure Interoperability: Tools and
Concepts* The Telecommunications Review, Mitretek Systems,
December 2002. See
http://www.mitretek.org/publications/2002_telecomm_review/stillson_07.pdf
- [PKIX] Stephen Kent, Tim Polk (co-chairs) *Public-Key Infrastructure IETF
Working Group*. See <http://www.ietf.org/html.charters/pkix-charter.html>
- [RFC3280] R. Housley, W. Polk, W. Ford, D. Solo *Internet X.509 Public Key
Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
IETF Network Working Group. April 2002
<http://www.faqs.org/rfcs/rfc3280.html>
- [Sun1] R. J. Perlman, Sun Microsystems, Inc *System and method for
shortening certificate chains* U.S. Patent Office, application
#20020147905 October 10, 2002. See
<http://appft1.uspto.gov/netahhtml/PTO/srchnum.html> (#20020147905)
- [X509] ITU-T (formerly CCITT). *Recommendation X.509: The Directory—
Public-Key and Attribute Certificate Frameworks*. 2000.

Dynamic Bridge Concept


Ken Stillson
Mitretek Systems

For presentation to the 3rd Annual PKI R&D Workshop

Problem Background

- Path discovery scaling problem
 - During work on path processing for the Federal Bridge CA, indications arose that even in the simplified mesh structure of a bridge, discovery can be very expensive
 - Time needed for a discovery of a new path increases much faster than linearly with the complexity of the mesh
 - Big problem when we reach big meshes (e.g. tying bridges together)

Path Discovery – The Sense of Direction Problem




- A “wrong-turn,” for example entering cloud X or Y, can lead to a wild goose chase during path discovery
- Path discovery has no “sense of direction,” no routing protocol, to determine a good route between sender and trust anchor.

Solution Background

- Some solutions noted to date
 - Pre-caching of objects (a “certificate spider”)
 - Removes network retrieval time and iterative discovery sequence; objects immediately available.
“Some assembly required.”
 - Pre-caching of partial paths
 - Attempt to cache common path segments; complicated by non-local constraints
(must be cached discovery only, not validation)


Dynamic Bridge Concept

- Pre-cache all possible paths from trust anchor to issuers (a “trust spider”), store “path minus 1” cached paths as direct cross-certificates from a new CA
 - The new cross certificates reflect the same trust arrangements as the initial mesh, but they have been “flattened” to combine intermediate hops




Dynamic Bridge Concept

- Dynamic bridge automatically creates Dyn(A) using a trust spider
- User of trust anchor A change their trust anchor to Dyn(A)
 - All paths to A now available from Dyn(A) with path length 1
- Walking a typical discovery algorithm:
 - Sender -> CA1 is easy; issuer in sender's cert, and often cert is included with message
 - Next step is to search for other certs with CA1's subject (searching for cross-certs issued to CA1). This immediately returns Dyn(A), which is known as the trust anchor, so discovery stops before starting.
 - No need to choose between alternative paths, so no direction needed




The Constraints Complication

- Consolidation needs to include intermediate constraints
 - In $A \rightarrow CA2 \rightarrow CA1$, the new constraint “ $Y \oplus X$ ” is the ordered combination of constraints Y and X .
 - E.g. if Y requires policy1, and X requires policy2, then $Y \oplus X$ would require both polices in a single constraint. If Y maps policy1 to policy2, and X maps policy2 to policy3, then $Y \oplus X$ map policy1 directly to policy3.
 - It is asserted that all the constraints possible in cross certificates can be combined by a fully implemented \oplus
 - “Last mile” validation still required for constraint Z



The Revocation Complication

- The dynamic bridge path has “short circuited” CA2
 - CA2 should be able to revoke it’s cross-cert to CA1, but dynamic bridge path does not pass through CA2, so removes revocation capability
- Solution: Set expiration date for the dynamic bridge’s certificate from Dyn(A) to CA1 to the earliest CRL “next update” for any CAs “flattened” from the path.
 - A dynamic bridge cross certificate is valid only until the time a CRL update could have invalidated part of the consolidated path
 - The dynamic bridge continuously re-generates as things expire. It re-performs validation, thus giving the opportunity to revoke.
 - Yes, the dynamic bridge will be “busy” in a complex mesh, but better a single bridge server per trust anchor than every desktop



Proposed Benefits


- Path discovery- iterative discovery is removed.
No wrong turns, as no turns
- Path validation- end-user validation is become simpler: path lengths are reduced and cumulative effects of the chained constraints are pre-calculated.
- Object location- the dyn-br has collected all objects a relying party will require. If dyn-br's cross-certs are all stored in a single directory, relying party need only point there. Dyn-br might also populate AIA fields in its cross-certs, to facilitate object location
- Construction requires no particular privilege, nor the cooperation of the mesh participants.
 - Any enterprise, even end-users, can establish their own
 - Those that utilize the dynamic bridge's trust anchor will receive the benefits, the existence of the dynamic bridge is non-harmful to those that do not utilize it.
- Utilization of a dyn-br does not require any specialized software
 - In-fact, only a subset of the standard path processing capabilities are needed; some PKI software that is not fully compliant now (due to lack of iterative path discovery, or limited object location capabilities) would be made fully capable

Unresolved Issues

- Presumably the dynamic bridge must use “on-line” keys
 - Common one-way push through firewall probably mitigates
- The concept hinges on the \oplus function, which is not yet implemented.
 - Is the X.509 constraint language sufficiently expressive to be associative?
- Microsoft Windows’s “CAPI” still has an object location challenge
 - No support for a “default directory” or an “AIA of last resort,” which could point to the dynamic bridge’s directory.
 - CAPI builds from the sender towards the trust anchor, so addition of AIA fields to dynamic bridge certificates does not help
- The spider crawl must start at the known trust anchor(s), and spread outwards. This direction is “the hard way” for object location
 - AIA is the “wrong direction.” SIA fields can be used for this, but are not widely populated. Bi-directional trust resolves this problem (pass 1 finds the “wrong-way” cert, then pass 2 knows the DN for the right-way). But uni-directional trust chops the search tree.

Conclusion

- Mitretek Systems is presenting the dynamic bridge concept to the PKI community without intent to assert patent protection, in hopes that its utility may be assessed, and discussions started concerning the possibility of implementation.
- Those interested in providing feedback, or joining discussions on the topic, are asked to contact the paper's author, Ken Stillson, at stillson@mitretek.org. If there is sufficient interest, a mailing list or similar discussion mechanism will be established.



GEMINI
SECURITY SOLUTIONS

Approaches to Certificate Path Discovery

Steve Hanna
Sun Microsystems

Peter Hesse
Gemini Security Solutions

Matt Cooper
Orion Security Solutions

Ken Stillson
Mitretek Systems

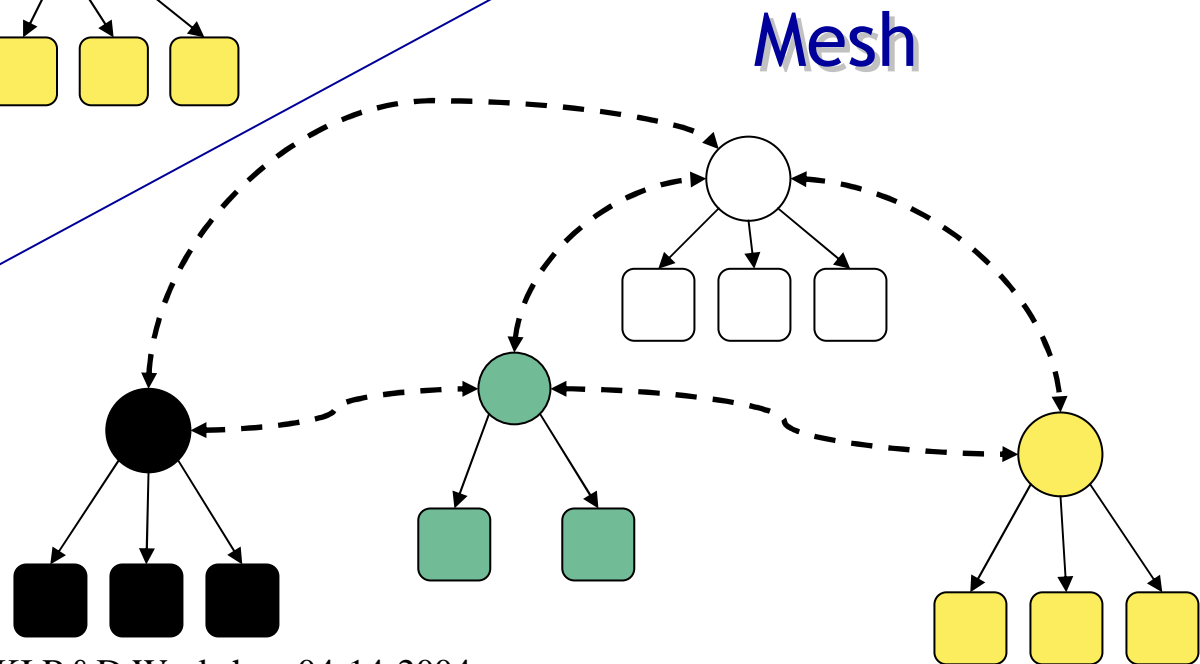

Agenda

- ▼ **PKI Structures**
- ▼ **Overview of Path Discovery**
- ▼ **Path Discovery Implementations**
 - Briefed by each panel member
- ▼ **Questions**
 - Prepared questions
 - Audience questions

PKI Structures


- ✓ **PKIs can be grouped into conceptual structures**
 - Hierarchical
 - Mesh
 - Bi-lateral Cross-Certified (Hybrid)
 - Bridge
- ✓ **Applications have commonly been developed assuming a particular structure**
 - This limits interoperability

Basic PKI Structures




Cross-Certified PKI Structures


Hierarchical




Hybrid



Mesh

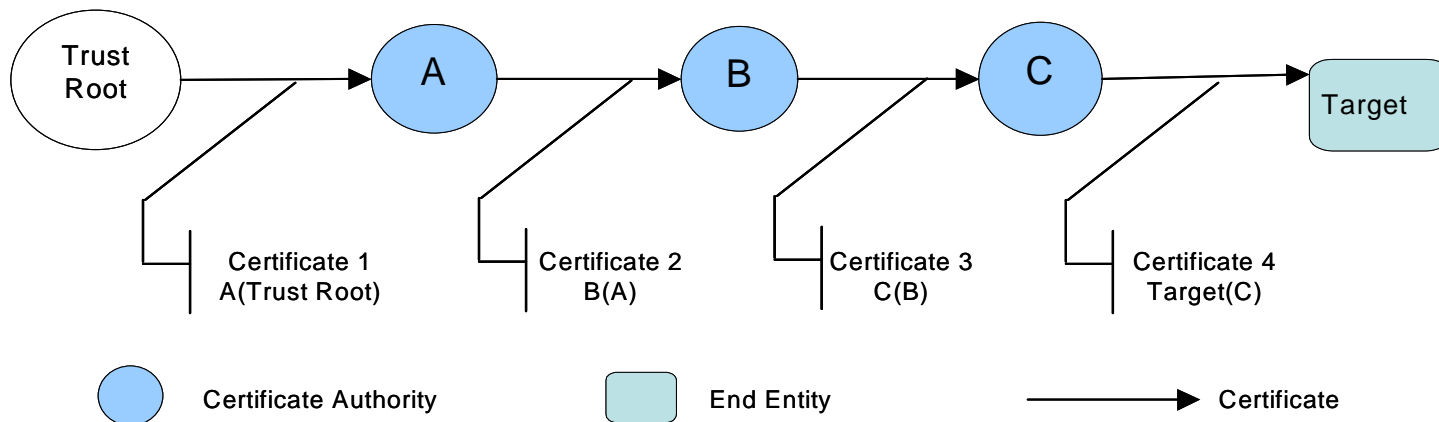


Bridge PKI Structure



Certification Path Building

- A certification path is an ordered list of certificates starting with a certificate issued by the relying party's trust root, and ending with the target certificate that needs to be validated



Certification Path Building (cont.)

- ✔ **Certification path building is not addressed by the standards that define the semantics and structure of a PKI**
 - Internet Draft in the works not as a standard but as an informational draft
- ✔ **The ability to construct or build a valid certification path is of paramount importance for applications that rely on a PKI**
- ✔ **Absent valid certification paths, you are working with PK—not PKI**

Path Building Implementations

- ✔ **Simply put, path-building is nothing more than a tree traversal**
 - certificates do not repeat in a path
- ✔ **Leads us to two basic algorithms:**
 - Start at a root, and work toward the end entity
 - Start at the end entity, and work toward a root
- ✔ **And, two models of implementation:**
 - Client-side
 - Server-side

Path Building Implementations

- ▼ Steve Hanna, Sun Microsystems
- ▼ Matt Cooper, Orion Security
- ▼ Ken Stillson, Mitretek Systems

Questions

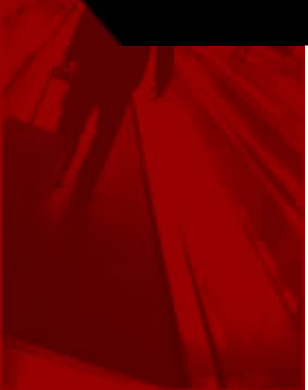

- ✓ **First, some prepared questions for our panelists:**
 - What is your implementation's goal when discovering paths? How do you attempt to reach that goal?
 - How does your implementation handle situations involving multiple trust roots?
 - How does your implementation reduce repetition of tasks between path discovery and path validation?

Questions (2)

- What are your recommendations for PKI architects/designers based upon path discovery?
- Give us ideas of your experiences with real/test PKIs (size, complexity, time to get a result, etc)
- Is path discovery best done on the client or server?
- What are the “next steps” in the world of path discovery?

Questions

▼ From the audience...



Johnson & Johnson

Use of Public Key Technology

Rich Guida

Director, Information Security

Johnson & Johnson

- **The world's largest and most comprehensive manufacturer of health care products**
- **Founded in 1886**
- **Headquartered in New Brunswick, NJ**
- **Sales of \$41.9 billion in 2003**
- **198 operating companies in 54 countries**
- **Over 110,000 employees worldwide**
- **Customers in over 175 countries**

Four Business Groups

- **Pharmaceuticals**
 - Prescription drugs including EPREX, REMICADE
- **Medical Devices and Diagnostics**
 - Blood analyzers, stents, wound closure, prosthetics, minimally invasive surgical equipment
- **Consumer Products**
 - E.g., Neutrogena; SPLENDA
- **Consumer Pharmaceuticals and Nutritionals**
 - E.g., TYLENOL

Statistics

- 400+ UNIX servers; 1900+ WinNT/2000 servers
- 96,000+ desktops/laptops (Win2K)
- 60,000+ remote users
 - Employ two-factor authentication (almost all using PKI; a few using SecurID but being migrated)
- 50M+ e-mails/month; 50+ TB of storage
- 530+ internet and intranet servers, 3.3M+ website hits/day

Enterprise Directory

- **Uses Active Directory forest**
 - Separate from Win2K OS AD but some contents replicated
- **Populated by authoritative sources only**
- **Uses World Wide Identifiers (WWIDs) as index**
- **Supports entire security framework**
 - Source of all information put into certificates
- **250K+ entries (employees, partners, retirees, former)**
- **LDAP accessible**

J&J PKI

- **Directory centric – certificate subscriber must be in Enterprise Directory**
- **Certificate contents dictated by ED info (none based on “user-supplied input”)**
- **Certificates issued with supervisor ID proofing**
- **Simple hierarchy – root CA and subordinate online CA**
- **Standard form factor: hardware tokens (USB)**
- **Production deployment began early 2003**
 - Total of over 105,000 certificates (signature and encryption) issued to date
- **Most important initial applications:**
 - Remote authentication
 - Secure e-mail
 - Some enterprise applications

Experience (1)

- **Training help desks (you can't do too much of this...)**
- **Ensuring sufficient help desk resources to respond to peaks (>100% of average level; fortunately reasonably short half-life)**
- **Shifting user paradigms (always hard to change human behavior...)**
 - **Patience**
 - **Clear, unequivocal instructions/steps**

Experience (2)

- **Hardware tokens**
 - CSP issues of “passphrase caching”
 - User recovery from lost, stolen or destroyed token
 - Short term recovery (network userID/PW)
 - Long term recovery (new cert(s))
- **Certificate revocation**
 - Reason codes in CRL (25% increase in size of CRL)
 - Don’t give users options to select (too confusing to them) – ask questions instead (then automate reason code selection)

Experience (3)

- **We put in three identifiers in each cert (e-mail address, WWID, UPN)**
 - **Right thing to do for apps**
 - **Means employee transfer out/transfer in processes require getting new certs (since e-mail address changes)**
 - **HR controls those processes, not IM**
 - **Moral: smart IM technical/policy decisions may require implementation outside IM**

Experience (4)

- **Once user gets new certs:**
 - Register them with apps (e.g., Outlook S/MIME profile changes)
 - Link them to other user accounts (e.g., Nortel VPN client)
- **Thus – there are some additional steps to “migrate” to new certs**
 - Not yet seamless

Experience (5)

- **Decryption private key recovery**
 - User can do for his/her own (after authenticating)
 - Local Key Recovery Authority Officer can request for others
 - Global KRAO must approve
 - But – important to distinguish key recovery from revocation or getting new certs
 - Unclear terminology (to users) resulted in lots of unnecessary requests, none of which required approval

Experience (6)

- **CRL growth is always faster than you predict**
 - Ours is approaching 1MB (expected it to be less than half that size)
- **Caching CRLs in Windows is “easy” but not obvious**
 - IE manages CRL cache as part of “temporary internet files” folder
 - Standard setting for us was: flush that folder when IE is closed
 - Results in lots of CRL downloads

Identifying and Overcoming Obstacles to PKI Deployment and Usage

Stephen R. Hanna
Sun Microsystems, Inc.
steve.hanna@sun.com

Jean Pawluk
Inovant
jpawluk@visa.com

Abstract

Public Key Infrastructure (PKI) is a fundamental security technology used in many applications. Nevertheless, PKI deployment has been slow. Why? In June and August 2003, the OASIS PKI Technical Committee conducted two surveys aimed at identifying the top obstacles to PKI deployment and usage and soliciting suggestions for how these obstacles can be overcome. This paper presents the results of those surveys and summarizes the PKI Action Plan that the PKI TC has developed in response.

1. Introduction

Around the world, security threats are escalating and the demands that business and personal information be safeguarded are mounting. Business, governments, and consumers want access to their information in a mode that is easy to use, yet secure.

Public Key Infrastructure (PKI) is a fundamental security technology used in many applications to provide those security assurances. For a number of years, the promise of PKI has been challenged by its complexity and the costs of deployment.

The OASIS PKI Technical Committee was formed in January 2003 to tackle the issue of how to successfully deploy and use Public Key Infrastructure. As early adopters of PKI technology, many members of the committee have first-hand experience with the challenges of implementing PKI technology. As a result of their combined experiences, the committee decided that an impartial survey was needed to further identify the critical obstacles to widespread use of PKI.

A short, multiple-choice web-based survey was prepared and hosted on the group's web site in June 2003. Invitations to participate in the survey were distributed to standards and industry groups as well as security vendors and their customers around the globe.

After reviewing the June 2003 survey results [1], the OASIS PKI Technical Committee prepared a second survey to gather more detailed data about specific obstacles. This second survey was publicized to the participants in the original survey during August 2003 [2].

The data gathered through these surveys provides a clear view of the obstacles impeding PKI deployment and usage. The survey respondents also provided specific suggestions for addressing these obstacles with a clear consensus emerging from the many responses.

Based on this consensus, the OASIS PKI Technical Committee developed a PKI Action Plan [3] with five specific action items addressing the top five obstacles identified in the surveys. After several months of public review and comment, the committee has published the PKI Action Plan and begun implementation.

Implementing the plan will require cooperation from many parties: vendors, customers, standards groups, etc. If these groups can overcome their differences and work together, the obstacles to PKI deployment may be greatly reduced.

2. Review of Previous Work

For several years, starting in 1997, the "Year of PKI" was proclaimed by vendors selling the promise that public key infrastructure would revolutionize security by safeguarding electronic transactions. While PKI has been very successful in certain realms (secure web browsing), the full scope of these declarations is yet to be fulfilled.

According to the findings of Burton Group research originally published in 2001 and in late 2002 [4], progress in PKI deployment has been made over the past decade, but very slowly. "While public key security potential is vast, public key infrastructure (PKI) continues to struggle with interoperability,

complexity and application integration issues that slow customer adoption. PKI's sophistication hasn't translated into mass enterprise deployments."

The Burton Group researchers state that "The major applications using PKI today remain web-based authentication and virtual private networks (VPN), though the use of digital signature based electronic forms applications continues to grow".

They also stated, "Much of the complexity retarding PKI arises because a complete PKI requires multiple products from multiple vendors" including the PKI enabled application, a certificate authority vendor, a directory services vendor, and the vendor specific software for hardware clients and servers. A functional PKI may also include scenarios "that include smart cards and other cryptographic devices, professional services or system integration services, access management portals, certificate validation services... and more".

These concerns about PKI are reflected in numerous similar articles and papers in the trade press, conferences, and workshops [5], [6].

3. June 2003 Survey Results

In the June 2003 survey conducted by the OASIS PKI Technical Committee [1], the participants were asked to rate the importance of several common PKI applications and the importance of commonly cited obstacles to PKI deployment and usage. They were also asked to provide demographic information, which was used to check for survey bias and correlations between demographics and opinions. Finally, they were asked to list applications and obstacles missing from the survey.

3.1. Survey Sample

The June 2003 survey was open to anyone with an opinion on PKI obstacles, but aimed at people with

expertise or experience in this area. Therefore, the survey invitations were sent to organizations and email discussion lists dedicated to PKI.

The 216 survey respondents were found to be a group of experienced group of industry professionals with serious PKI experience.

A large variety of job titles and functions were found among the respondents. Many of them had both technical and business functions included within their scope of their job duties. More than 75% of the respondents had at least 5 years of experience in Information Security / Privacy.

With over 90% of the respondents having either deployed or developed PKI software, they were very experienced with PKI. The majority of the participants were from the USA and Canada (60%) however over 30 countries were represented with many participants from Europe or Asia.

3.2. Analysis of Applications

Survey respondents were asked to rate various PKI applications as Most Important, Important, or Not Important to them. Respondents were also able to enter their own application area under Other (such as Identity Management, Non-Repudiation, and Document Encryption) and rate its importance.

For analysis, these ratings were combined into a weight by assigning 2 points for each respondent who rated an application Most Important and 1 point for each rating of Important. By computing these weights, the applications can be ranked by importance (as indicated by the respondents).

As shown in Table 1, most applications were found to be important but no one application stood out as the most important.

Applications	Most Important	Important	Not Important	No Answer	Weight	Weight Rank
Document Signing	43%	47%	6%	3%	1.38	1
Web Server Security	42%	48%	6%	4%	1.37	2
Secure Email	40%	46%	8%	6%	1.33	3
Web Services Security	34%	53%	9%	4%	1.26	4
Virtual Private Network	33%	50%	11%	6%	1.24	5
Electronic Commerce	34%	48%	13%	5%	1.22	6
Single Sign On	28%	56%	12%	4%	1.17	7
Secure Wireless LAN	25%	48%	19%	8%	1.06	8
Code Signing	20%	50%	22%	8%	0.98	9
Secure RPC	6%	40%	40%	13%	0.61	10
Other Application	9%	3%	7%	81%	0.21	11

Table 1: Application Weight Rank

Application weights are shown graphically in Figure 1.

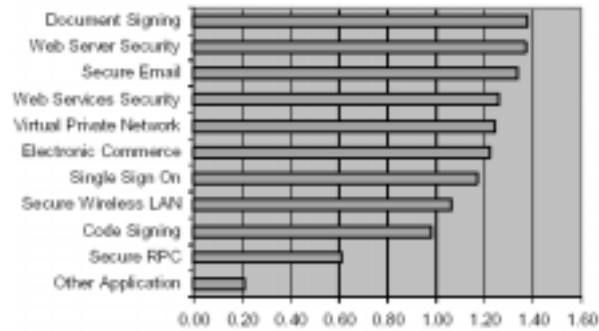


Figure 1 PKI Application Weights

These results affirm the view that PKI is a foundational technology used in many applications.

As business, governments and consumers all have different PKI needs, they also have different concerns about the importance of the listed applications. The survey results showed strong correlations between respondents' employment sector and their rating of applications. Government sector respondents ranked Document Signing 10% higher and Code Signing 11% lower than the total sample. In contrast, respondents in the Computer-related Manufacturing sector ranked Code Signing 12% higher than the total sample and Document Signing 10% lower. This is not surprising, since governments produce a lot more documents than code and computer firms typically do the opposite.

3.3. Analysis of Obstacles

In a manner similar to the rating of applications, respondents were presented with a list of possible obstacles to PKI deployment and usage and asked to rank each one as a Major Obstacle, a Minor Obstacle, or Not an Obstacle. Respondents were also able to describe an obstacle under Other and rate it in the same way.

Weights were computed by assigning 2 points to Major Obstacles and 1 point to Minor Obstacles. Using these weights, ranks were computed. The results are shown in Table 2.

The PKI Obstacles weight ranking is shown graphically in Figure 2.




Figure 2 PKI Obstacle Weights

Many survey participants listed other obstacles to PKI deployment and usage. Here is a list of the obstacles that were cited by several respondents:

- Insufficient ROI/business justification/need
- Enrollment too complicated
- Smart card problems (cost, driver and OS problems, readers rare)
- Revocation hard
- Standards (too many, incompatible, changing, poorly coordinated)
- Too much focus on PKI technology, not enough on business need
- No universal CA
- Too complex
- Insufficient skilled personnel
- Poor implementations

Obstacles	Major Obstacle	Minor Obstacle	Not an Obstacle	No Answer	Total	Weight	Weight Rank
Software Applications Don't Support It	54%	33%	10%	3%	100%	1.45	1
Costs Too High	53%	34%	12%	2%	100%	1.42	2
PKI Poorly Understood	47%	41%	11%	1%	100%	1.37	3
Poor Interoperability	46%	39%	12%	3%	100%	1.35	4
Hard to Get Started - Too Complex	46%	39%	13%	2%	100%	1.34	5
Hard for End Users to Use	43%	42%	13%	3%	100%	1.30	6
Lack of Management Support	30%	44%	21%	5%	100%	1.09	7
Too Much Legal Work Required	25%	50%	22%	3%	100%	1.03	8
Hard for IT to Maintain	20%	55%	21%	4%	100%	0.99	9
Other Obstacle	18%	3%	5%	74%	100%	0.39	10

Table 2: PKI Obstacles Weight Rank

Unfortunately, the outcome of the survey question on obstacle ratings was inconclusive. Many obstacles had similar weights. Obstacles were broadly defined so it was not clear what respondents meant. In addition, several obstacles cited as Other Obstacles were noted by multiple respondents, indicating that the list of obstacles was incomplete. Therefore, the OASIS PKI Technical Committee Survey decided to conduct a followup survey to clarify the obstacles and ratings.

4. August 2003 Survey Results

The OASIS PKI Technical Committee’s August 2003 survey [2] introduced a new points-based rating system that allowed respondents to clearly indicate priorities. It added “Other” obstacles cited by multiple participants in the June 2003 survey. It asked several questions designed to refine the broad categories used in the June survey. Moreover, it asked respondents to suggest ways that the obstacles could be addressed.

4.1. Survey Sample

The OASIS PKI Technical Committee sent invitations only to people who responded to the June 2003 Survey and provided an email address. This allowed us to use the previously gathered demographic data in analyzing the results while avoiding the need to ask for such data again. We found that the respondents to the August 2003 survey were similar in demographics and opinions to the earlier respondents.

4.2. Analysis of Obstacles

Instead of asking respondents to rate obstacles as a Major Obstacle, a Minor Obstacle, or Not an Obstacle, the August 2003 survey asked respondents to allocate 10 points among the obstacles listed, giving points to each item according to its importance. This allowed respondents to heavily weight items that were especially important to them. The results are shown in Table 3.

Obstacle	Average Points	Rank
Software Applications Don't Support It	1.76	1
Costs Too High	1.26	2
PKI Poorly Understood	1.06	3
Too Much Focus on Technology, Not Enough On Need	1.01	4
Poor Interoperability	.90	5
Hard to Get Started – Too Complex	.68	6
Lack of Management Support	.66	7
Hard for End Users to Use	.59	8
Enrollment Too Complicated	.35	9
Too Much Legal Work Required	.33	10
Smart Card Problems	.32	11
Hard for IT to Maintain	.30	12
Insufficient Need	.29	13
Revocation Hard	.25	14
Standards Problems	.25	15

Table 3: PKI Obstacles Point Rank

The point-based rankings reveal a substantial difference between the top five obstacles, which account for about 60% of the points, and the remaining ten obstacles. This does not mean that the lower-rated obstacles are not important. Most of them were rated as Most Important or Important by a majority of the respondents to the June 2003 survey. But the top five obstacles are just *more* important to the survey respondents.

The results were carefully checked for any sign that a small number of respondents might be skewing the results by throwing more votes than average to one item. This was not found to be true. In fact, the obstacle rankings were consistent across many demographic lines (experience, geography, industry sector, etc.). This was true for almost all opinions expressed in both surveys (except application ranking, as noted above).

Perhaps the most valuable part of the Follow-up Survey was the textual responses. For each of the top obstacles identified in the June 2003 Survey, respondents were asked to describe in their own words what causes these obstacles and what the PKI TC or others could do to address the obstacles. Certain themes were repeated over and over by many respondents. These themes pertain to several of the top obstacles. They are:

- Support for PKI is inconsistent. Often, it’s missing from applications and operating systems. When present, it differs widely in what’s supported. This increases cost and complexity substantially and makes interoperability a nightmare.
- Current PKI standards are inadequate. In some cases (as with certificate management), there are too many standards. In others (as with smart cards), there are too few. When present, the standards are too flexible and too complex. Because the standards are so

flexible and complex, implementations from different vendors rarely interoperate.

5. PKI Action Plan

The two surveys conducted in June and August 2003 allowed the OASIS PKI Technical Committee to identify the primary obstacles to PKI deployment and usage and to develop a PKI Action Plan [3] to address the obstacles. Here is a brief synopsis of that Action Plan.

5.1. Call for Industry-Wide Participation

The OASIS PKI Technical Committee recognizes that it cannot act independently in implementing this Action Plan. PKI involves many parties: customers and users, CA operators, software developers (for applications, PKI components, platforms, and libraries), industry and standards groups, lawyers, auditors, security experts, etc. This PKI Action Plan was developed based on input from all of these parties. The OASIS PKI Technical Committee calls on these parties to assist in its implementation.

5.2. Action Items

Develop Application Guidelines for PKI Use

For the three most popular PKI applications (Document Signing, Secure Email, and Electronic Commerce), specific guidelines should be developed describing how the standards should be used for this application. These guidelines should be simple and clear enough that if vendors and customers implement them properly, PKI interoperability can be achieved.

PKI TC members will contact application vendors, industry groups, and standards groups to determine whether such guidelines already exist and if not who could/should work on creating them. In some cases, standards may need to be created, merged or improved. If application guidelines already exist, the PKI TC will simply point them out.

Who: PKI TC Guidelines Subcommittee, Application Vendors, and Industry and Standards Groups

When: Spring 2004 for initial work

Increase Testing to Improve Interoperability

Provide conformance test suites, interoperability tests, and testing events for the three most popular applications (Document Signing, Secure Email, and Electronic Commerce) to improve interoperability.

Certificate management protocols and smart card compatibility are also a concern. Branding and certification may be desirable. The PKI TC will work with organizations that have demonstrated involvement in or conduct of PKI interoperability testing or conformance testing to identify and encourage existing or new efforts in this area. Interoperability has many aspects. See the PKI Interoperability Framework white paper at <http://www.pkiforum.org/whitepapers.html> for details.

Who: PKI TC Testing Subcommittee with Industry and Standards Groups

When: Spring 2004 for initial work

Ask Application Vendors What They Need

OASIS PKI TC members will ask application vendors for the three most popular applications (Document Signing, Secure Email, and Electronic Commerce) to tell us what they need to provide better PKI support. Then we will explore how these needs (e.g. for quantified customer demand or good support libraries) can be met.

Who: PKI TC Ask Vendors Subcommittee, in cooperation with application vendors

When: Spring 2004 for initial work

Gather and Supplement Educational Materials on PKI

Explain in non-technical terms the benefits, value, ROI, and risk management effects of PKI. Include specific examples of PKI applications with real benefits and ROI. Also explain when PKI is appropriate (or not). Educational materials should be unbiased and freely available to all. If these materials already exist, the PKI TC will simply point them out. Otherwise, it will develop them in cooperation with others.

When: January – August 2004

Explore Ways to Lower Costs

Encourage the software development community (including the open source community) to provide options for organizations to conduct small pilots and tests of PKI functionality at reasonable costs—in effect reducing cost as a barrier to the use of PKI. Of course, operating a production PKI involves many costs other than software acquisition so an effort will be undertaken to gather and disseminate best practices for cost reduction in PKI deployments around the world.

Who: PKI TC Lower Costs Subcommittee, software development community, customers, etc.

When: Initial efforts in 2004

6. Conclusions

The results of the surveys conducted by the OASIS PKI Technical Committee identify the primary obstacles to PKI deployment and usage, as judged by the survey respondents. They also provide suggestions for addressing those obstacles.

Based on these results and on feedback from many PKI users, vendors, and other stakeholders, the OASIS PKI Technical Committee has prepared a PKI Action Plan to address the obstacles identified. Implementing the PKI Action Plan will be challenging but it provides some hope that PKI deployment will be easier and the benefits of PKI (strong and scalable security) will be widely realized.

7. Acknowledgments

Without the hard work and dedication of the members of the OASIS PKI Technical Committee, the results documented here would never have come to light. The survey respondents are thanked for their hard-won insights, which serve as a primary source for this paper. In addition, the PKI Action Plan reviewers and supporters are thanked for their assistance in hopes that it will lead to the successful completion of the PKI Action Plan.

Thanks to OASIS for granting permission for portions of the PKI Action Plan and survey analyses to be reproduced in this paper.

Thanks to The Burton Group for allowing us to quote one of their research reports.

8. References

[1] OASIS Public Key Infrastructure Technical Committee, "Analysis of June 2003 Survey on Obstacles to PKI Deployment and Usage", August 8, 2003
<http://www.oasis-open.org/committees/pki/pkiobstaclesjune2003surveyreport.pdf>

[2] OASIS Public Key Infrastructure Technical Committee, "Analysis of August 2003 Follow-up Survey on Obstacles to PKI Deployment and Usage", October 1, 2003
<http://www.oasis-open.org/committees/pki/pkiobstaclesaugust2003surveyreport.pdf>

[3] OASIS Public Key Infrastructure Technical Committee, "PKI Action Plan", February 2004
<http://www.oasis-open.org/committees/pki/pkiactionplan.pdf>

[4] Dan Blum and Gerry Gebel, "Public Key Infrastructure: Making Progress, But Many Challenges Remain, V2", Directory and Security Strategies Research Report No. 612, Burton Group, Utah, February 13, 2003, pp. 5-7, <http://www.burtongroup.com>

[5] United States General Accounting Office, "Status of Federal Public Key Infrastructure Activities at Major Federal Departments and Agencies", (GAO-04-157), Washington D.C., December 2003
<http://www.gao.gov/cgi-bin/getrpt?GAO-04-157>

[6] Peter Gutmann, "PKI: It's Not Dead, Just Resting", IEEE Computer, August 2002, pp. 41-49

9. Copyright Notices

Copyright 2004 Sun Microsystems, Inc. All Rights Reserved.

Copyright 2004 Inovant. All Rights Reserved.

For portions reproduced from OASIS documents:

Copyright OASIS Open 2004. All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to OASIS, except as needed for the purpose of developing OASIS specifications, in which case the procedures for copyrights defined in the OASIS Intellectual Property Rights document must be followed, or as required to translate it into languages other than English.

OASIS PKI TC:

Identifying and Overcoming Obstacles to PKI Deployment and Usage


Jean Pawluk (Inovant) & Steve Hanna (Sun)

April 2004

Acknowledgements

OASIS Public Key Infrastructure Technical Committee -

A dedicated group of PKI technology “early adopters” including



D E S I G N
D E V E L O P
D E P L O Y

Assumptions

- Public Key Infrastructure (PKI) is a fundamental security technology
- PKI's promise as a foundation technology is challenged by its very complexity & the costs of deployment.

OASIS PKI Technical Committee was formed in January 2003 to tackle the issue of how to successfully deploy and use Public Key Infrastructure


The Surveys

If PKI is such a useful technology why isn't more widely used ?

PKI TC wanted more objective viewpoints:

Two surveys commissioned:

- June 2003 - Initial Survey
- August 2003 - Detailed Survey




The Approach

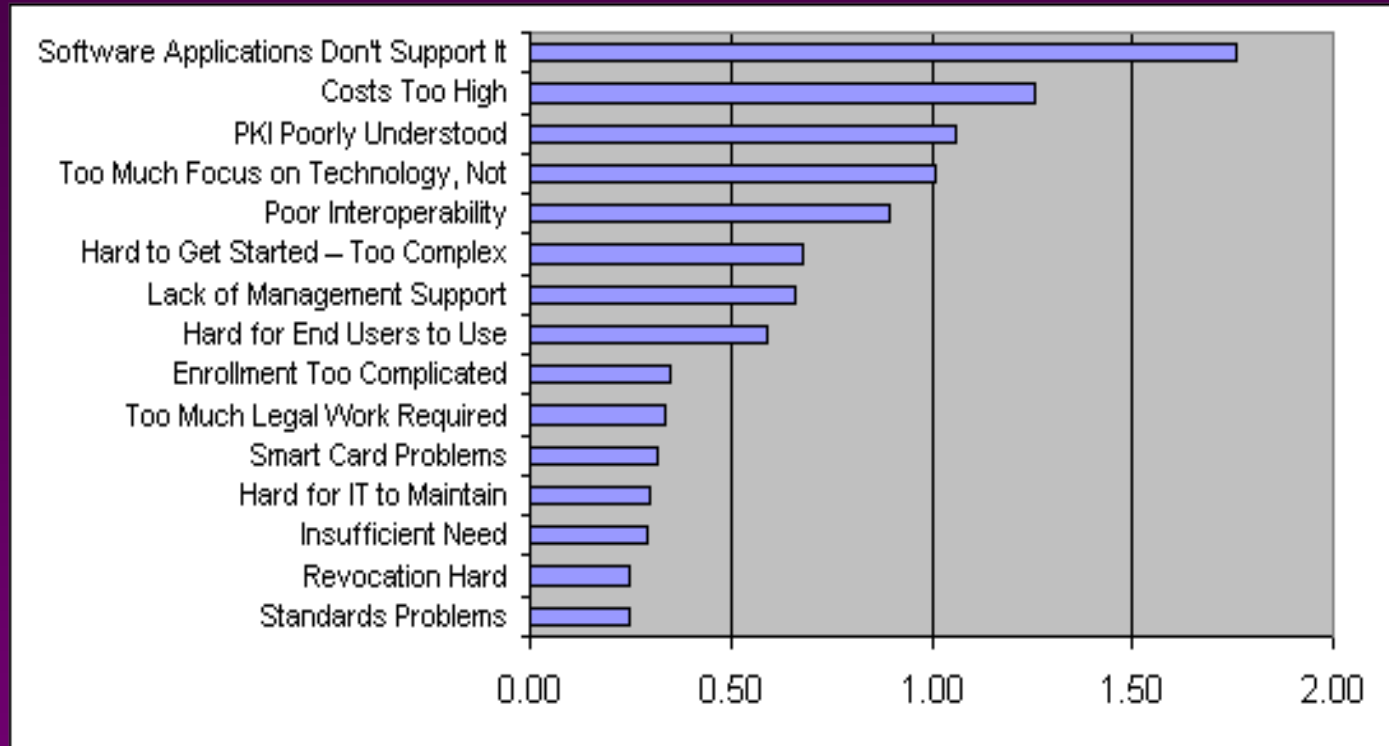
Survey invitations sent to organizations and email discussion lists dedicated to PKI.

The 216 survey respondents are a group of experienced group of industry professionals with serious PKI experience.

- Over 90% of the respondents have either deployed or developed PKI software

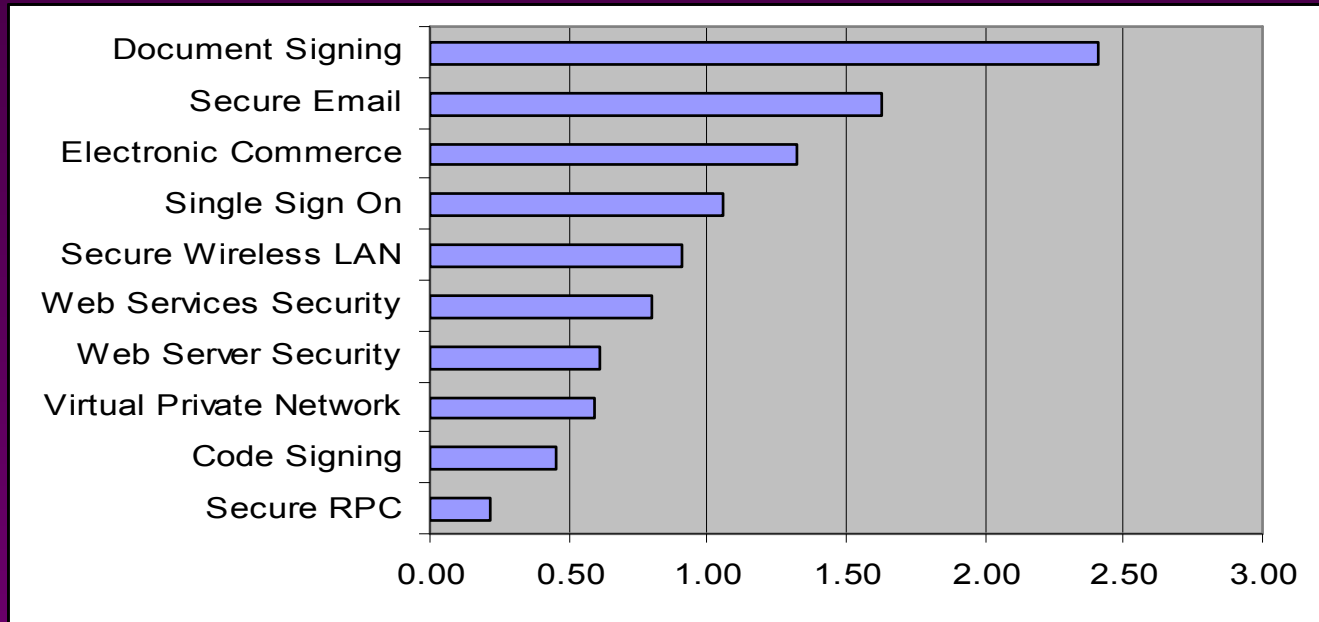


Obstacles: Ranked by Importance



The first four obstacles have more than half of the total points


Applications: Ranked by Need for Improvements in PKI Support




Support for PKI is inconsistent.

- Often, it's missing from applications and operating systems or if present, it differs widely in what's supported.
- Current PKI standards are inadequate
 - In some area (as with certificate management there are too many standards. In others (e.g. smart cards), there are too few

Costs Ranked by Most Problematic




Parties: Ranked by Greatest Need for PKI Understanding



Few understand what is the value of PKI

Where the Most Serious Interoperability Problems Arise



Frustration level with PKI results from attempts to implement and having serious interoperability problems

PKI Call to Action - 1

Develop Application Guidelines for PKI Use

Create specific guidelines for three most popular PKI applications describing how the standards should be used for this application.

- Document Signing,
- Secure Email
- Electronic Commerce

These guidelines should be simple and clear enough that if vendors and customers implement them properly, PKI interoperability can be achieved.

PKI Call to Action - 2


Increase Testing to Improve Interoperability

- Provide conformance test suites, interoperability tests, and testing events for the three most popular applications
 - Document Signing
 - Secure Email
 - Electronic Commerce
- Certificate management protocols and smart card compatibility are a concern.
- Branding and certification may be desirable.

PKI Call to Action - 3

Ask Application Vendors What They Need


- Ask application vendors to tell us what they need to provide better PKI support.
- Explore how these needs (e.g. for quantified customer demand or good support libraries) can be met.



PKI Call to Action - 4

Gather and Supplement Educational Materials on PKI

- Explain in non-technical terms the benefits, value, ROI, and risk management effects of PKI.
- Include specific examples of PKI applications with real benefits and ROI.
- Explain when PKI is appropriate (or not).



PKI Call to Action - 5

Explore Ways to Lower Costs

Reduce cost as a barrier to the use of PKI.

- Encourage the software development community (including the open source community) to provide options for organizations to conduct small pilots & tests of PKI at reasonable cost.


Operating production PKI involves many costs other than software acquisition

- Gather “best practices” for cost reduction in PKI deployments.

Join Us ...

OASIS Public Key Infrastructure Technical Committee has begun implementation of its PKI Action Plan

<http://www.oasis-open.org/committees/pki/pkiactionplan.pdf>



D E S I G N
D E V E L O P
D E P L O Y

End Users Viewpoint


Who do you trust ?

Legal Contracts & Assumed Risk

Liability Issues

- **Identity Binding**

- **Cross Chaining vs. Closed Systems Validation**



PKI Obstacles and Action Plan (My top three list)

Sean W. Smith

**Department of Computer Science
Dartmouth College**

`www.cs.dartmouth.edu/~sws/`

April 14, 2004



Missing Obstacle #3

Does it say what you want it to say?

Suppose we closed our wireless network with EAP-TLS.

- Why can't end users authorize visitors?
- (Why did we have to play tricks with SPKI/SDSI in cookies?)
- Why can't the Dartmouth net recognize a Princeton visitor?

Why are proxy certificates necessary?

Why will the doctor's office have a post-it note with the PIN?

Action: Find better ways to have signed assertions follow real-world trust flow

Missing Obstacle #2

Do the humans understand it?

Can Johnny encrypt yet?

Is it easy to do the right thing?

Do mental models match what the machines are doing?

Action: HCISEC.

Missing Obstacle #1

Does it work?

PKI is a lot of work. But there's a point to it.

Besides asking...

- “Do the users get it?”
- “Does the code work?”

Action:...we should also ask:

- “What are the security goals of using PKI in this application, and do we achieve them?”

Server-side SSL?

Question: If Alice's browser gives her all the right signals that she has an SSL connection to Bob's server, does she?

What we learned: with Netscape/Linux and IE/Windows then current, no. With a lot of work, you can add a trusted path to Mozilla.

Ye Smith 2002

<http://www.cs.dartmouth.edu/~sws/abstracts/ys02.shtml>

Digital Signatures?

Question: Does Bob's valid digital signature on document D mean that Bob approved the contents of D ?

What we learned: With standard office tools and many "best of breed" PKI tools, it was easy to construct documents:

- whose contents changed in usefully malicious ways
- without invalidating the signature

Kain Smith Asokan 2002

<http://www.cs.dartmouth.edu/~sws/abstracts/ksa.shtml>

Client-side SSL?

Question: If Alice submits her request to Bob's server with client-side SSL under her cert, did Alice approve that request?

What we learned:

- If the adversary has access to a server, be careful
- With IE, if the adversary gets a user-level program on your machine, game's over...
- ...even with hardware tokens

Marchesini Smith Zhao 2003, 2004

<http://www.cs.dartmouth.edu/~sws/abstracts/msz04.shtml>