Royal Holloway University of London

Information Security Group

New architectures for identity management: Unifying security infrastructures

Chris Mitchell
Royal Holloway, University of London

Royal Holloway University of London

Information Security Group

Acknowledgements

- This is joint work with Haitham Al-Sinani, also of Royal Holloway.

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

Royal Holloway University of London

Information Security Group

User authentication

- The need for authentication of human users is a fundamental security requirement (perhaps ***the*** fundamental requirement).
- Particularly relevant to EuroPKI, since PKI is all about supporting human authentication.
- Despite its importance, it is almost universally acknowledged that providing user authentication remains a huge practical problem.

Royal Holloway University of London

Information Security Group

Passwords

- In practice, as many observers have noted, we are still using passwords for almost everything.
- Again, as widely acknowledged, the use of passwords has many shortcomings, not least because users today have so many Internet relationships, all needing authentication.
- In such a context, password re-use and use of weak passwords are almost inevitable.

5

Royal Holloway University of London

Information Security Group

Solutions

- The usual approach to this problem is to propose yet another new way of achieving user authentication, possible involving a PKI.
- However, perhaps there are already enough good technological solutions?
- Maybe the problem is adoption of the solutions we already have? How do we fix this?
- Of course, this is partly a business case/ sociological issue, but maybe it is also a problem which requires new technical thinking?

6

Royal Holloway University of London

Information Security Group

New thinking required

- It is easy for those of us doing technical research to claim that this is not our problem.
- We provide the technology and the business/ commercial world should just get on with it.
- However, life is not so simple.
- We as academics should be thinking about how to devise technological solutions which are easier to adopt.
- Key issues for easy adoption are transparency, ease of use, and backwards compatibility.

7

Royal Holloway University of London

Information Security Group

Identity management

- Identity management systems have been designed to simplify user authentication.
- Such a system enables an Identity Provider (IdP) to support authentication of a User (and assertion of user attributes) to a Service Provider (SP).
- Recent years have seen the emergence of a wide range of such systems, e.g. OpenID, Liberty, Shibboleth, CardSpace and OAuth.
- Each has its own set of protocols governing communications between the main parties.

8

Royal Holloway University of London

Information Security Group

Infrastructure support

- As well as its own protocols, each system may also have a unique supporting infrastructure, including public key certificates, shared keys, passwords, etc.
- Some systems have gained traction recently, e.g. Facebook's adoption of OAuth (Facebook Connect), and significant use of OpenID.
- However, the systems that have been most widely used are also those which have the most significant problems (e.g. phishing vulnerabilities).

9

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

10

Royal Holloway University of London

Information Security Group

Security infrastructures

- In order to use cryptography to protect communications, some kind of security infrastructure needs to be in place.
- In its simplest form, this will just be a means to set up shared secret keys between communicating parties.
- Traditionally, e.g. in banking networks, this can be achieved using one or more Trusted Third Parties (TTPs).
- One type of TTP for this purpose is known as a Key Distribution Centre (KDC).
- A KDC shares a secret key with every party, and these keys can be leveraged (using an appropriate protocol) to set up a secret key between any two parties.

11

Royal Holloway University of London

Information Security Group

Public Key Infrastructures (PKIs)

- As we all know, a PKI is simply another type of security infrastructure, based on digital signatures.
- A Certification Authority (CA) creates digitally signed certificates for user public keys, binding a user name to a public key.
- If universally adopted, PKIs could provide a robust underpinning for user authentication.

12

The promise of a universal PKI

- Some years ago, PKI was the subject of huge hype.
- Companies producing PKI products (e.g. CA software) or providing PKI services suddenly (and temporarily!) became hugely valuable.
- In many cases the vision sold as part of this hype was of some kind of universal PKI, whereby every PC in the world would have a public key certificate, which could then be used for a huge range of purposes, e.g.:
  - secure e-commerce;
  - universal secure e-government;
  - secure home banking;
  - electronic signatures for all;
  - …

PKI – what happens in practice I

- Of course, this has not happened.
- There are many PKIs, each set up for a specific purpose.
- For example:
  - companies have their own PKIs, used to support internal secure communications;
  - MasterCard and Visa (and card issuing banks) have PKIs set up to support EMV (used to support smart card based credit/debit card transactions, e.g. in parts of Europe);
  - Internet web sites have certificates used for SSL/TLS security.
- There are, of course, many explanations for this – one being the fact that the policies under which certificates are issued will depend on the context of use.

Royal Holloway University of London

Information Security Group

PKI – what happens in practice II

- More generally, PC users do not have the expertise or motivation to generate a signature key pair, and obtain a certificate for their public key.
- This can be seen from the failure of the SET e-commerce secure payment system; a major obstacle to its adoption was the need for every user to generate a key pair and take a copy of their public key to their bank.
- End users cannot be expected to understand the operation of public key cryptography.
- Moreover, current PCs often do not have a means for secure key storage (needed for the private key), although TPMs may help.

15

Royal Holloway University of London

Information Security Group

An evolutionary approach

- One major problem with introducing a large scale PKI is the huge cost and complexity, and issues such as the need for user education.
- The scheme we propose is designed to enable an evolutionary adoption of more secure means of authentication, avoiding the need for a 'big bang'.
- Unless there is a very strong business case, such changes are very hard to engineer (see, for example, problems with introducing ID cards in the UK).

16

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

17

Royal Holloway University of London

Information Security Group

Well known problems

- We start by reviewing some of the well known problems with existing authentication solutions.
- These problems apply very broadly.

18

Royal Holloway University of London

Information Security Group

The phishing threat

- Many identity management systems are susceptible to phishing attacks, in which a malicious (or fake) SP redirects a user browser to a fake IdP.
- The user then reveals to the fake IdP secrets that are shared with a genuine IdP.
- This arises because, in the absence of a system-aware client agent, schemes rely on browser redirects.

19

Royal Holloway University of London

Information Security Group

Lack of consistency

- One huge problem faced by any user is that the user experience of every identity management system is different.
- We all know that users fail to make good security decisions, even when confronted with relatively simple decisions.
- The lack of consistency is likely to make the situation much worse, with users simply not understanding the complex privacy- and security-relevant decisions they are being asked to make.

20

Privacy

- When using third party IdPs which provide assertions about user attributes, there is a danger that a user will damage their privacy by revealing attributes unintentionally to an SP.
- This is a threat when using systems like OAuth (e.g. as instantiated by Facebook Connect).
- In general, getting privacy settings right is highly non-trivial.

Another new infrastructure?

- It is tempting to try to devise another new scheme which has the practical advantages of OAuth and OpenID, but yet provides robust protection against phishing and privacy loss.
- That is, devise a client-based scheme with the user convenience of other systems, but which somehow avoids the fate of CardSpace.

Royal Holloway University of London

Information Security Group

Problems

- However, it seems that a new solution is:
  - unlikely to succeed when others (some with a great deal of inertia and incorporating very nice features, e.g. CardSpace) have failed;
  - likely to create yet another different user experience, increasing the likelihood of serious mistakes.
- Thus maybe this is not the right approach.

23

Royal Holloway University of London

Information Security Group

A new approach?

- The goal of this talk is to consider a new approach to the problem.
- It does not involve proposing any new protocols or infrastructures.
- The goal is to try to make it easier to use existing systems, and also to make their use more secure (less prone to phishing) and privacy-enhancing (consistent interface and explicit consent).

24

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

25

Royal Holloway University of London

Information Security Group

Client-based solution

- The scheme we propose involves a client-based user agent.
- This is a single tool which supports a wide range of ID management systems yet provides a single interface to the user.
- The consistent user interface should maximise user understanding of what is happening (and reduce risk of errors).
- It also avoids the need for passive browser redirects, hence mitigating phishing attacks.

26

Royal Holloway University of London

Information Security Group

Motivation for scheme

- One motivation for the scheme comes from considering CardSpace (and its open source 'twin', Higgins).
- Before proceeding we thus need to spend a bit of time describing CardSpace.

27

Royal Holloway University of London

Information Security Group

CardSpace: a brief description

- CardSpace acts as client-based agent, and provides a consistent card-based user interface.
- That is, sets of user credentials (relationships with IdPs) are represented to users as cards.
- CardSpace also defines a set of protocols for interactions between IdPs, Clients (user machines) and SPs.

28

CardSpace operation

- The user, interacting with the browser via the *identity selector*, may have identities issued by one or more IdPs.
- Each identity is represented by an *InfoCard* held by the identity selector, and this InfoCard is the means by which the user interacts with the identity selector to choose which identity to use.
- Each IdP runs a Security Token Service (STS), to generate security tokens.
- A *Self-issued Identity Provider* may be provided by a client platform to allow use of self-issued tokens.

CardSpace Identity Selector

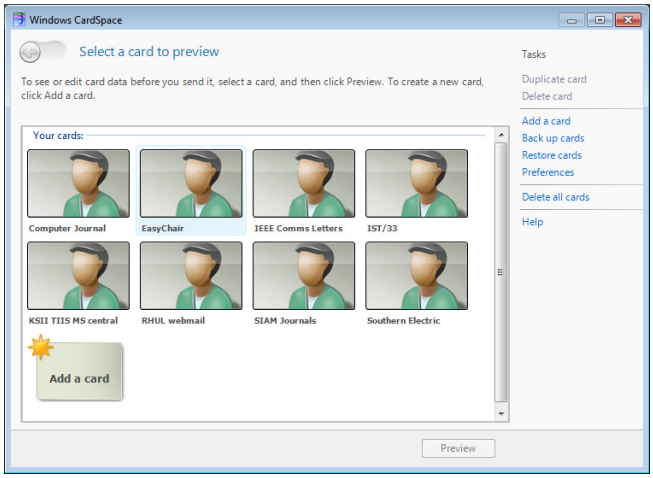

Royal Holloway University of London

Information Security Group

CardSpace architecture

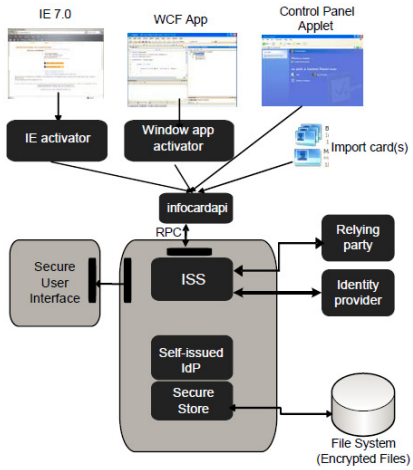

31

Royal Holloway University of London

Information Security Group

CardSpace interaction model

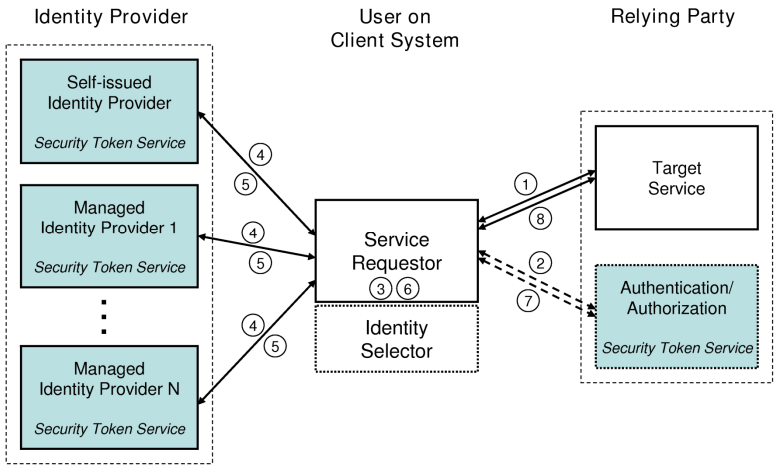

32

Royal Holloway University of London

Information Security Group

Operation I

1. Service requester gets the security policy of the target service. We suppose that the policy requires the requester to get a token issued by an IdP's STS.
2. (optional) The service requester gets the policy of the authentication/authorisation STS (to determine properties of required token).
3. The requester asks the identity selector to provide a security token meeting the policy of the target service.
4. The identity selector first gets the user to choose an InfoCard capable of meeting the target service requirements, and then gets the policy of the selected IdP's STS.

33

Royal Holloway University of London

Information Security Group

Operation II

5. The InfoCard indicates the method to be used to authenticate the user to the IdP STS; the user sends an appropriate credential to the IdP STS, and the identity selector gets back a token.
6. The token is given to the service requester.
7. (optional) The service requester presents the token to the STS, which generates a token for the target service.
8. The service requester presents the token to the target service to get access.

34

Royal Holloway University of London

Information Security Group

User authentication

- Before issuing a token, an IdP will typically need to authenticate the user.
- This user authentication takes place via the local CardSpace software
- Two key advantages:
  - provides consistent user experience;
  - limits possibility of phishing attacks.

35

Royal Holloway University of London

Information Security Group

An observation

- The user interface of CardSpace and the underlying communications protocols are not inherently tied together.
- Why not keep the simple/intuitive user interface, and use it as the front end for a tool which manages user credentials in a consistent way regardless of the underlying identity management system?

36

Royal Holloway University of London

Information Security Group

An observation (contionued)

- Credential sets could identify with which identity management system (or systems) they should be used.
- For example, each credential set could be stored as a self-describing XML document.
- Indeed, these credential sets could include username/password pairs.

37

Royal Holloway University of London

Information Security Group

A universal client adapter

- We can now describe our scheme.
- We call it *IDSpace* (homage to the role CardSpace played in developing our idea).
- IDSpace has two main components – a browser plugin (the *IDSpace extension*) and a separate piece of software (the *IDSpace client software*).
- Both execute on the user platform.

38

Royal Holloway University of London

Information Security Group

IDSpace high level architecture

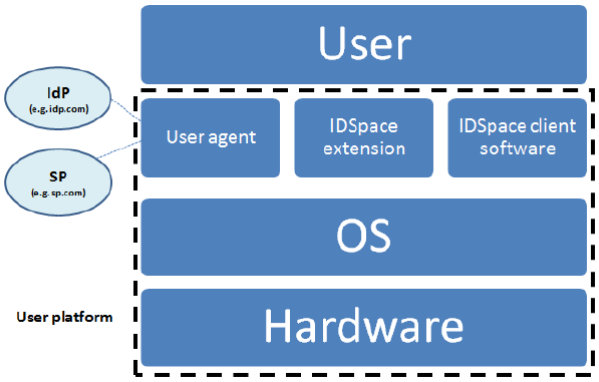

39

Royal Holloway University of London

Information Security Group

IDSpace components

- The IDSpace system possesses a number of components, as shown on the next slide:
  - **Card Selector**: presents a card-based interface to user to enable choice of IdP and credentials;
  - **cCard store**: stores credential cards (cCards) containing credential info (used by Card Selector);
  - **Credential store**: separate secure storage for keys, passwords, attributes, etc., associated with cCards;
  - **Kernel**: core component controlling system operation;
  - **Page Scanner**: scans web pages;
  - **Activator**: activates the Card Selector.

40

Royal Holloway University of London

Information Security Group

IDSpace client architecture

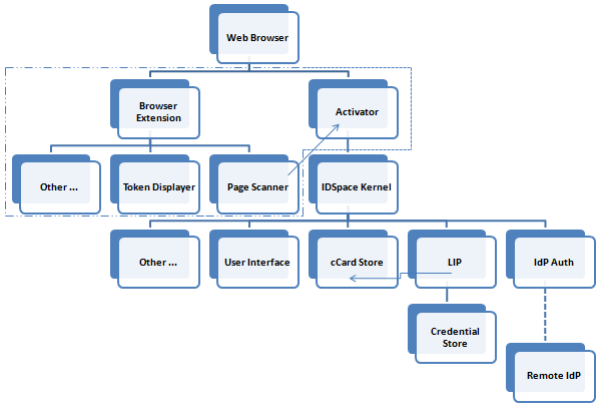

41

Royal Holloway University of London

Information Security Group

Sketch of protocol I

- The IDSpace works as follows.
  1. User browses an SP login page.
  2. The *IDSpace Page Scanner* examines the page to see which identity systems are supported.
  3. The user is offered a choice (e.g. via right clicking) of systems to use. [There are many options for implementing this step.]
  4. The IDSpace Activator activates the *IDSpace Card Selector*.

42

Sketch of protocol II

5. The *IDSpace Data Transporter* passes metadata (e.g. selected identity system, SP identity, SP policy) to the *IDSpace Kernel*.
6. The *Kernel* interacts with the *Card Selector*, which allows the user to choose a cCard (and possibly an identity system).
7. The *Kernel* interacts with the selected IdP to obtain a token for use by the SP. If necessary the IdP authenticates the user via the *Card Selector*.
8. The *Token Displayer* asks the user for permission to send the token to the SP.

User experience

- The user interacts with a single piece of software (the Card Selector) regardless of which underlying system is in use.
- This enables the user to use a single simple interface to:
  - choose (and manage) credentials;
  - be authenticated to an IdP;
  - give consent for release of PII to an SP.

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

Royal Holloway University of London

Information Security Group

Mappings

- Each identity system operates differently, and hence each system maps slightly differently onto IDSpace.
- Main relevant characteristic is whether an identity system is:
  - redirect-based, or
  - active-client-based.
- We look at these two cases.

Active-client-based systems

- Here a browser incorporates an 'active client', which acts as an intermediary between SPs and IdPs, and is aware of the identity system.
- All SP-IdP communications involve this active client.
- The active client might prompt the user to select a digital identity, choose an IdP, review an identity token created by the IdP, and/or approve a transaction.
- Phishing attacks are mitigated.
- The active client can also give a consistent user experience and a greater degree of user control.
- Examples include CardSpace and Liberty (when using a Liberty-enabled client (LEC)).

Role of IDSpace (active client case)

- In such a case the IDSpace client software plays the role of the active client.
- IDSpace acts as a type of 'universal' client, integrating the various systems and handling credential information and user authentication in a unified and consistent way.
- Thus, for example, IDSpace can transparently replace the Microsoft CardSpace software.

Royal Holloway University of London

Information Security Group

Redirect-based systems

- In such a scheme, the browser is redirected by an SP to an IdP (and vice versa).
- Typically, such schemes work on unmodified browsers.
- A major disadvantage is that a malicious SP can redirect the browser to a fake IdP (e.g. to fraudulently obtain user credentials).
- Examples include OpenID, Liberty (browser-post profile), Shibboleth, and Facebook Connect (OAuth).

49

Royal Holloway University of London

Information Security Group

Role of IDSpace (redirect case)

- The IDSpace client software essentially converts a redirect system into an active-client system.
- Redirects are no longer under the control of the SP (and IdP).
- The IDSpace client also manages authentication of the user to the IdP.
- The operation of IDSpace is completely transparent to the IdP and SP.

50

Royal Holloway University of London

Information Security Group

Features

- Regardless of the ID system protocols supported by the SP and IdP, IDSpace is transparent to both parties.
- That is, no parties (except the user who installs and uses the software) need to be aware of its presence.
- As long as the SP and IdP share at least one identity system, then IDSpace operation is possible.

51

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

52

Royal Holloway University of London

Information Security Group

Password management

- Password managers are commonplace.
- However, apart from schemes built into browsers, they do not appear to be widely used.
- *PassCard* is a browser-plugin-based scheme we have described previously which allows CardSpace to be used as a password manager.
- The idea behind PassCard could readily be extended to allow IDSpace to provide password management facilities (with username/password pairs being represented as cCards).

53

Royal Holloway University of London

Information Security Group

Moving into the cloud

- Cloud-based identity management systems offer some advantages over client-based schemes (not least portability).
- Indeed, cloud-based variants of CardSpace have been proposed in which InfoCards are cloud-based.
- One possible extension of IDSpace would be to make it cloud-based.

54

Royal Holloway University of London

Information Security Group

Identity system interoperation

- In other work, we have proposed and prototyped a series of client-based (browser plug-in based) schemes to support interoperation between an IdP and an SP supporting different identity systems.
- This functionality could also be supported by an IDSpace client.

55

Royal Holloway University of London

Information Security Group

Role of client agents

- IDSpace is just one example of the potential power of a client-based security agent.
- There are many other ways in which client-hosted software might be used to assist users in making difficult security-relevant decisions when using Internet services.
- Indeed, this paper is really intended to encourage the research community to think more about using client-based schemes to improve user security.

56

Royal Holloway University of London

Information Security Group

Agenda

- Introduction
- Security infrastructures
- A problem – not another infrastructure
- A solution – not another infrastructure!
- Mappings to specific systems
- Other functionality
- Concluding remarks

57

Royal Holloway University of London

Information Security Group

IDSpace works!

- A preliminary prototype of IDSpace has recently been built by my co-author (Haitham Al-Sinani), and is still under development.
- Unfortunately it is not yet in a demonstrable state.
- However, we soon hope to make available a usable prototype.

58

Royal Holloway University of London

Information Security Group

Related work

- Copies of published papers on PassCard and the various identity management interoperation schemes can be found on my home page:

  www.chrismitchell.net
- Many are also available as RHUL technical reports:

  www.ma.rhul.ac.uk/tech

59

Royal Holloway University of London

Information Security Group

Questions?

- For further information please contact:
  - Haitham Al-Sinani
    Haitham.Al-Sinani.2009@live.rhul.ac.uk
  - Chris Mitchell
    me@chrismitchell.net
- Address:

  Information Security Group
  Royal Holloway
  University of London
  Egham TW20 0EX
  UK

60