

Interoperation between Identity Management Systems

Chris Mitchell
Information Security Group
Royal Holloway, University of London

Agenda

1. Identity Management Systems
2. Current systems
3. Differences and issues
4. Interoperation
5. Privacy and security
6. Concluding remarks

Need for identity management

- Today's user typically has many accounts with many Internet service providers.
- Each account has its own 'name' for the user, and also its own *credential*, i.e. a means of authenticating the user (e.g. a password).
- This presents a huge burden on the user, who typically resorts to one or both of two bad practices:
 - writing down passwords, or
 - re-using passwords.

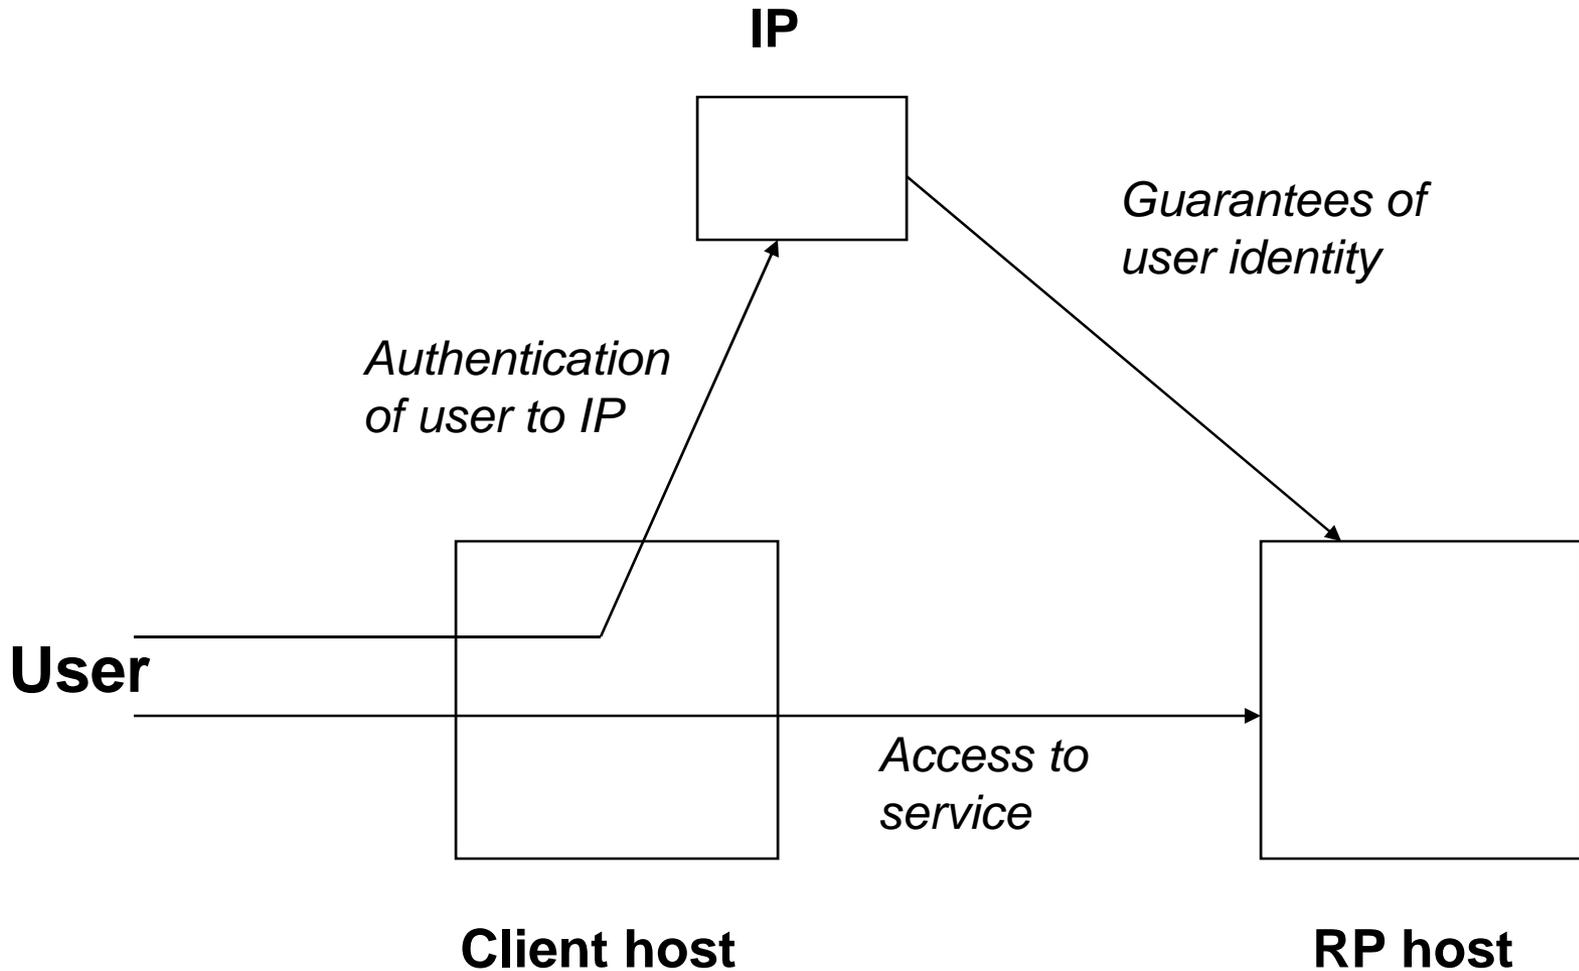
What is an ID Management System?

- Various definitions in use.
- For purposes of this talk ...
 - an Identity Management System (IDMS) is a scheme which enables a user to delegate some of the responsibility for credential management to a TTP called an *Identity Provider (IP)*;
 - this reduces task of credential management for user (at cost of delegating trust).

Roles

- **User** – human being for whom service is ultimately provided;
- **Client** – platform (e.g. PC) employed by User;
- **Relying Party (RP)** – provider of service which wants assurance about user identity;
- **Identity Provider (IP)** – authenticates user/client and then provides assurances about user to RP.

Operation



Terminology

- **Single Sign-On (SSO):** an SSO system is a special type of IDMS in which user authenticates to IP just once and then is authenticated automatically to multiple RPs.
- **User-centric:** a user-centric IDMS is simply an IDMS in the terminology used here.
- **Claim-based:** a claim-based IDMS is one in which the IP not only authenticates the user, but may store other information about the user (attributes).

Agenda

1. Identity Management Systems
2. Current systems
3. Differences and issues
4. Interoperation
5. Privacy and security
6. Concluding remarks

History

- An early example of a (failed) IDMS is provided by Microsoft Passport.
- Microsoft introduced Passport:
 - provided an SSO service for Passport-registered users to Passport-registered RPs;
 - no longer operates as an SSO service – used simply as a means of managing Microsoft logins.
- This seems to have acted as a spur to the industry, and there are now a whole range of IDMSs.

SAML overview

- SAML (Security Assertion Markup Language) is an OASIS standard.
- Actually two major versions: 1.1 and 2.0 (with significant differences).
- Standards specify two (quite different) things, both designed to support IDMSs:
 - SAML assertions – XML data structures;
 - protocols to support an IDMS.
- Arguably SAML is not actually an IDMS, but certainly provides key messaging components of an IDMS.

SAML assertions

- SAML assertions are a standardised means of enabling one party (e.g. an IP) to make statements about authentication of a User.
- Three types:
 - Authentication statements;
 - Attribute statements;
 - Authorisation decision statements.
- These standardised assertions are widely used in IDMSs.

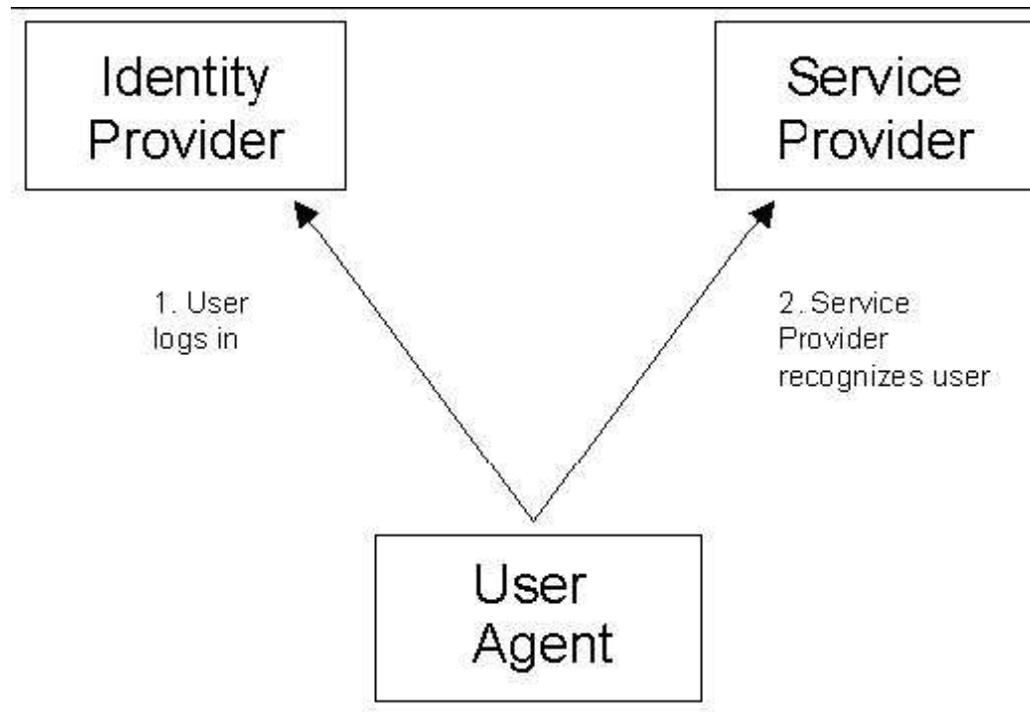
SAML protocols

- These enable SAML assertions to be transferred from an IP to an RP in response to a *query* from the RP.
- Three types of query:
 - Authentication queries;
 - Attribute queries;
 - Authorisation decision queries.
- SAML v2.0 also includes protocols for other functions necessary to support an IDMS (e.g. ‘single logout’).

Liberty Alliance

- The Liberty Alliance is a consortium of companies interested in SSO and identity management.
- It has published a series of specifications for an 'open' XML-based SSO system.

Liberty SSO Model



Role of Identity Provider

- In Liberty, a User authenticates to a Liberty *Identity Provider* (IP).
- The IP then automatically authenticates user to RPs.
- User then needs only one password (or other means to authenticate to IP).
- Works using http redirection.

Liberty operation I

- Typical operational scenario is as follows.
- User visits web site of RP, and SSL connection established.
- RP then redirects user web browser to Liberty IP which establishes SSL connection and then authenticates the user (if necessary).
- Liberty IP then redirects user back to RP.

Liberty operation II

- Messages need to be passed between RP and IP.
- RP sends authentication request and IP responds with authentication response (containing 'security assertions').
- Messages passed either embedded in URLs or in http forms (using POST method).
- Syntax of messages based on SAML.

Pseudonymity

- Liberty requires the IP to use a different pseudonym with each RP.
- Gives a level of unlinkability for users (a type of anonymity).
- However, may be compromised through network addresses.

CardSpace

- CardSpace is a Microsoft architecture for identity management.
- It has a number of component parts:
 - A distributed architecture for identity management;
 - A set of defined Web Services interfaces between entities in the architecture;
 - A set of software is available both for Windows Vista and XP which will enable users to manage their identities in a Windows environment;
 - Development support to enable applications to use CardSpace managed identities.

Identity Metasystem

- Microsoft refers to this collection of components as an Identity Metasystem.
- The idea is to provide a unified way for (Windows) users to use many different underlying identity management systems.
- Key ideas here are:
 - provide a simple user model for identity;
 - enable users to control which identity is used for what purpose through user interface notion of **InfoCards**.

Some CardSpace definitions

- *Digital identity*: a set of claims made by one digital subject about itself or another digital subject.
- *Digital subject*: a person or thing, represented or existing in the digital realm.
- *Claim*: an assertion of the truth of something.

CardSpace comments I

- The Microsoft definition of digital identity is a very general one, and does not distinguish between two concepts which are often treated separately:
 - identifiers or labels (e.g. email address, National Insurance Number, passport number, ...);
 - attributes (e.g. the identity holder is an employee of company X, a silver card holder for airline Y, a season ticket holder for train route Z, ...)

CardSpace comments II

- There are two main justifications for the Microsoft 'claims' approach:
 - it enables protocol interactions to be simplified – a single protocol can be used to transfer claims;
 - some types of claim are difficult to categorise – a credit card number may be viewed as both an identifier and an attribute.
- However, on the down side, human beings by and large understand the distinction between the two types of claim – this means that it may be a useful distinction.
- Thus CardSpace is a **claim-based** IDMS.

OpenID

- OpenID is a decentralised SSO system (with some similarities to Liberty) – it is open source.
- Users register with an OpenID identity provider (IdP).
- A service provider using OpenID displays a login form containing a space for an OpenID identifier, indicating a particular identity with a particular IdP (no password).

Using OpenID

- The RP then communicates with the appropriate IdP, either via the user's browser or directly.
- The user's browser is redirected to the IdP, and, if necessary the IdP then authenticates the user (OpenID does not restrict how this is done).
- The IdP then redirects the user's browser back to the RP and provides an authentication assertion.

Adoption and issues

- Use of OpenID is growing rapidly.
- The technology is now backed by a lot of leading players (Google, Microsoft, ...).
- See www.openid.net
- As with all systems relying on redirection at the behest of the RP, the scheme is open to phishing attacks if username/password used for authentication.

OpenID and CardSpace

- Because CardSpace and its identity metasystem are token-format-agnostic, CardSpace does not compete directly with other Internet identity architectures like OpenID.
- In some ways, OpenID and CardSpace can be seen as complementary.
- Indeed, CardSpace Information Cards can be used today for signing into OpenID providers, Windows Live ID accounts, SAML identity providers, and other services.

Shibboleth

- Shibboleth is an architecture and implementation for a federated identity-based authentication and authorisation system.
- Identities are treated as attributes, as in CardSpace.
- It is open source.
- Shibboleth builds on SAML 1.1.

Higgins

- Higgins is an open source identity framework with significant similarities to CardSpace.
- Like CardSpace is uses a card-based metaphor for managing user identities.

Agenda

1. Identity Management Systems
2. Current systems
3. Differences and issues
4. Interoperation
5. Privacy and security
6. Concluding remarks

Common features

- **Use of SAML:** all the schemes we have discussed either mandate SAML, or, in the case of the two frameworks, support its use.
- All adhere to general IP-RP model discussed earlier in talk.

Differences in scope

- CardSpace and Higgins are frameworks, i.e. in some sense are not complete IDMSs.
- They are designed to allow a variety of token types to be used to support IDM.
- Nevertheless, they do have standardised message formats/types.
- CardSpace also does not easily support SSO.
- Liberty, OpenID and Shibboleth, however, are complete schemes, with fixed token types.

Open-ness

- Shibboleth, Higgins and OpenID are open source, in that software is freely available,
- However, the term is not so relevant to Liberty, which is just a set of specifications – there could be open source Liberty implementations out there ...
- CardSpace is not open source, but nevertheless all major interfaces are public (except Windows interface).

Agenda

1. Identity Management Systems
2. Current systems
3. Differences and issues
4. Interoperation
5. Privacy and security
6. Concluding remarks

Need for interoperation

- Many systems are being deployed.
- If the user has a different user experience, and a different set of identity providers for almost every different service, then the whole point of IDMSs will be lost.
- Thus, ideally, either one system will win out, or some level of interoperation between systems is needed.

Who should support interoperation?

- Of course, if every RP supports every IDMS, then there will essentially be no problem.
- However, this seems unlikely (and would be a big burden on small service providers).
- If every IP supports multiple systems, then the problems will be less (but user will still have variety of UXs).
- Client-side support for interoperation would be useful.

Barriers to interoperation

- Differences in scope are a major issue for a client-side approach.
- If RP is using a claim-based IDMS, it may expect to interact with an IP that can make assertions about a range of user attributes.
- However, an IP supporting an ‘authentication only’ IDMS, e.g. Liberty, will not generate such assertions.

Concordia

- The Concordia project (www.projectconcordia.org) is a global initiative designed to drive interoperability across identity protocols in use today.
- It solicits and defines real-world use cases and requirements for use of multiple identity protocols in practical deployment scenarios.
- It supports the creation of protocol solutions.

Concordia – recent work

- Concordia demonstrated technology-provider interoperation of high-priority scenarios at the RSA conference in April 08.
- This involved InfoCard – federation, and SAML v2.0 – WS-Federation chaining scenarios.

A client-based approach

- A client-based approach to interoperation between CardSpace and Liberty was presented in:

W. Alrodhan and C. J. Mitchell, 'A client-side CardSpace-Liberty integration architecture', Proceedings of IDtrust 2008, 7th Symposium on Identity and Trust on the Internet, NIST, Gaithersburg, MD, March 2008.

Integrating the two schemes

- *Identity management architecture adaptor* is software installed on client which understands Liberty and CardSpace message flows/formats.
- Interposes itself between IPs and SPs adhering to different identity management architectures, to translate messages.
- In case of Liberty IP and CardSpace RP, we assume that there is a pre-established trust relationship (including pseudonyms, and an InfoCard identifier).

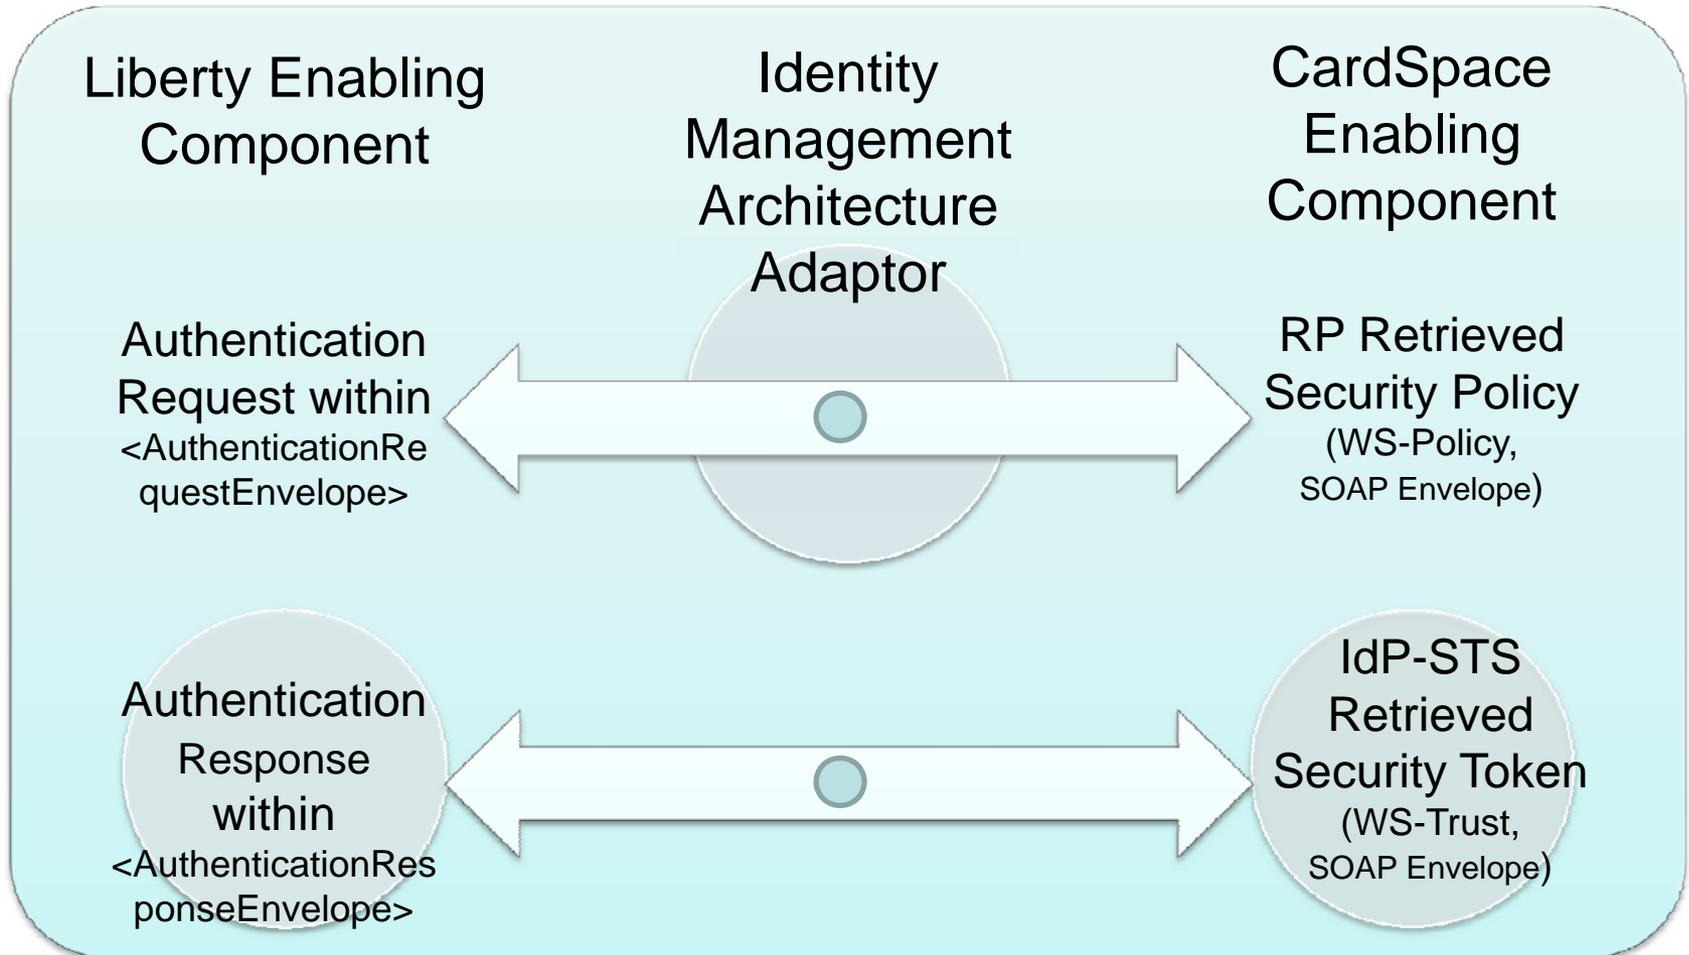
Restrictions

- The scheme has the following restrictions:
 - Only supports SAML tokens;
 - Only asymmetric proof of rightful possession of the token (holder-of-key);
 - In case of CardSpace RP & Liberty IP, token freshness requests are discarded.

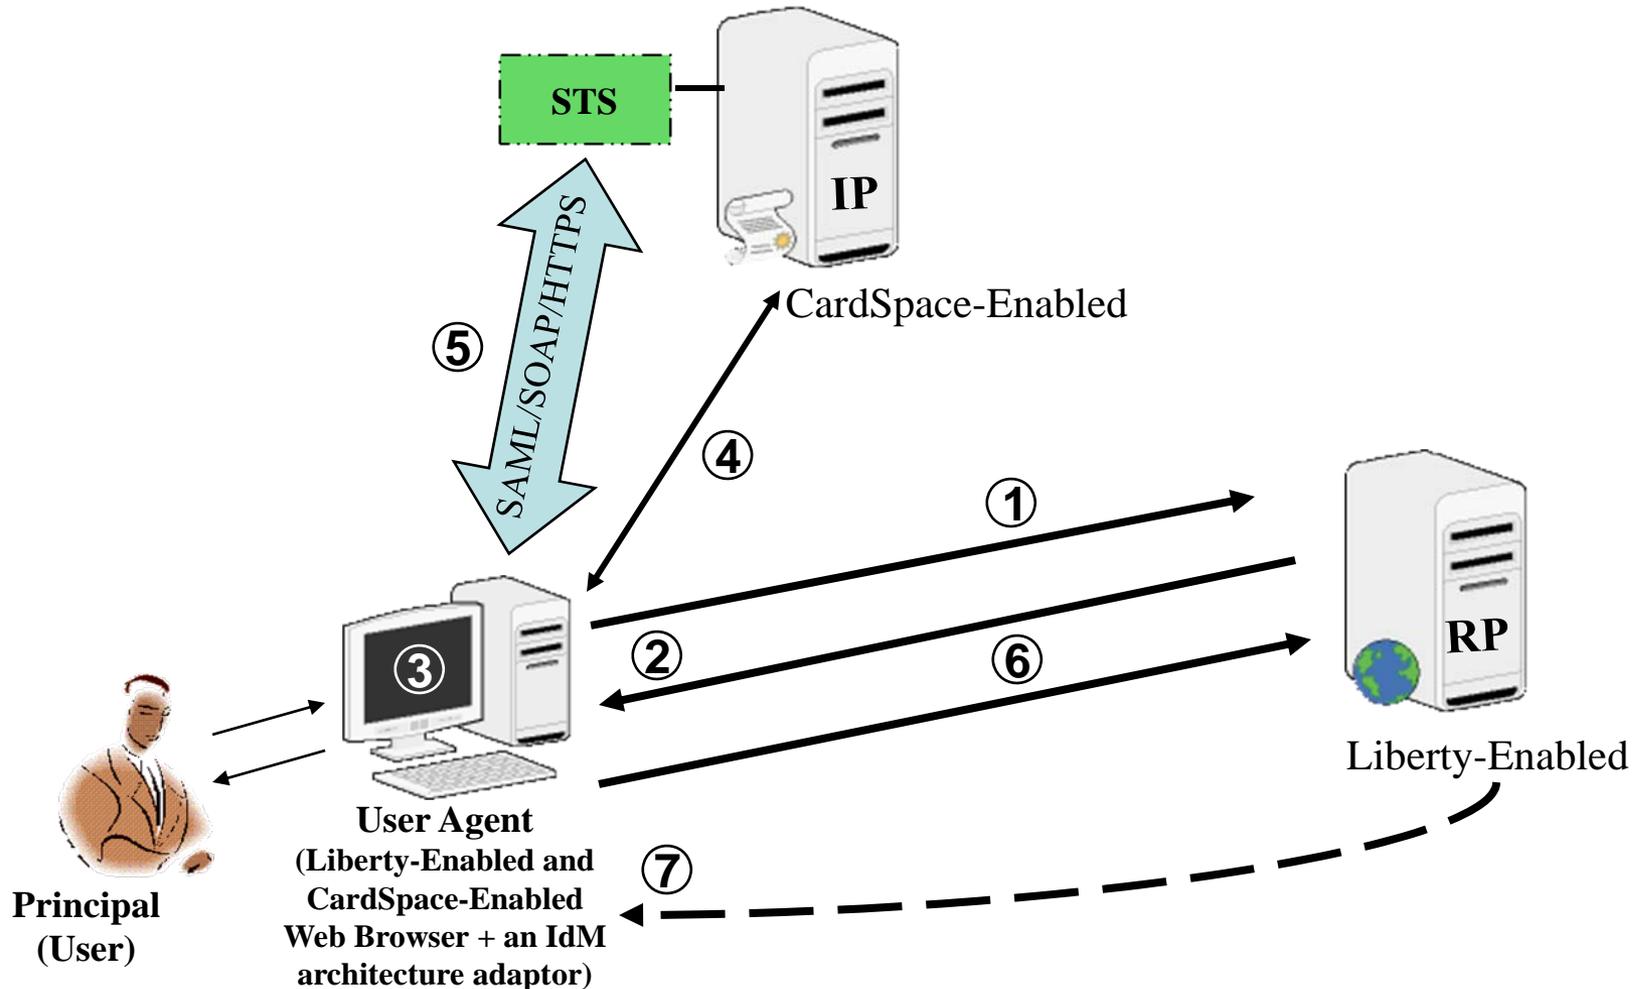
Representing the claims

- Possible approaches:
 - **SAML attribute statement:** Would require some modifications to the Liberty enabling component;
 - **Authentication with no claims:** Severe impact on the usability of the integrated scheme.

Data flows



A possible scenario



Analysis

- The integration model does not require Microsoft/Liberty cooperation.
- However, implementing such a model is non-trivial task.
- CardSpace and the Liberty ID-FF have somewhat different scopes.
- User-agents still need to be CardSpace and Liberty enabled.
- There is no end-to-end encryption

Agenda

1. Identity Management Systems
2. Current systems
3. Differences and issues
4. Interoperation
5. Privacy and security
6. Concluding remarks

Phishing attacks

- There is a major ‘phishing’ problem with any IDMS (how does the user know the IP is genuine?).
 - If the IP uses a password to authenticate the user, then compromise of this password is potentially very serious.
 - In a claim-based IDMS, i.e. where IP potentially holds PII about user, then need to be very careful about how this information is managed and disseminated.

Claim-based systems

- Certain privacy issues arising in CardSpace have been discussed in:
 - W. Alrodhan and C. J. Mitchell, 'Addressing privacy issues in CardSpace', in: *Proceedings of IAS '07, Third International Symposium on Information Assurance and Security, Manchester, UK, August 2007*, IEEE Computer Society (2007), pp.285-291.
- These issues are largely the same for any claim-based IDMS.
- (Above paper also proposes possible solutions).

CardSpace issues

- CardSpace, like many other IDMSs, has a number of limitations, including:
 - Reliance on DNS names to identify IPs and RPs;
 - In default CardSpace scenario, IP is aware of the identities of the RPs (to prevent token replay attacks using “symmetric” means);
 - Reliance on user’s judgment of RP trustworthiness (which gets user PII);
 - Reliance on a single layer of authentication.

RP trustworthiness

- User judgment regarding the honesty of the RP is a security-critical task.
- RP will obtain user's personal information in form of "asserted claims" within a security token.
- Within CardSpace, user judgment is based on one of:
 - RP's high-assurance public key certificate.
 - RP's 'ordinary' public key certificate (e.g. a certificate used for SSL/TLS);
 - No certificate at all.

Authentication

- Session security relies on a single layer of authentication, e.g. using an X.509 certificate, Kerberos v5 ticket, self-issued token or password.
- If working session is hijacked (e.g. by compromising a self-issued token), or password is cracked (e.g. via guessing, brute-force, key logging, or dictionary attacks), security of the whole system is compromised.

Agenda

1. Identity Management Systems
2. Current systems
3. Differences and issues
4. Interoperation
5. Privacy and security
6. Concluding remarks

Where next?

- Interoperation between IDMSs is a high priority issue – needs more research and more development/testing.
- Privacy and security issues inherent in IDMSs need to be addressed.
- In long term, to avoid high risk compromise of IP credentials, need to give users stronger credentials, e.g. card/token-based, and employing public key cryptography backed by a PKI.