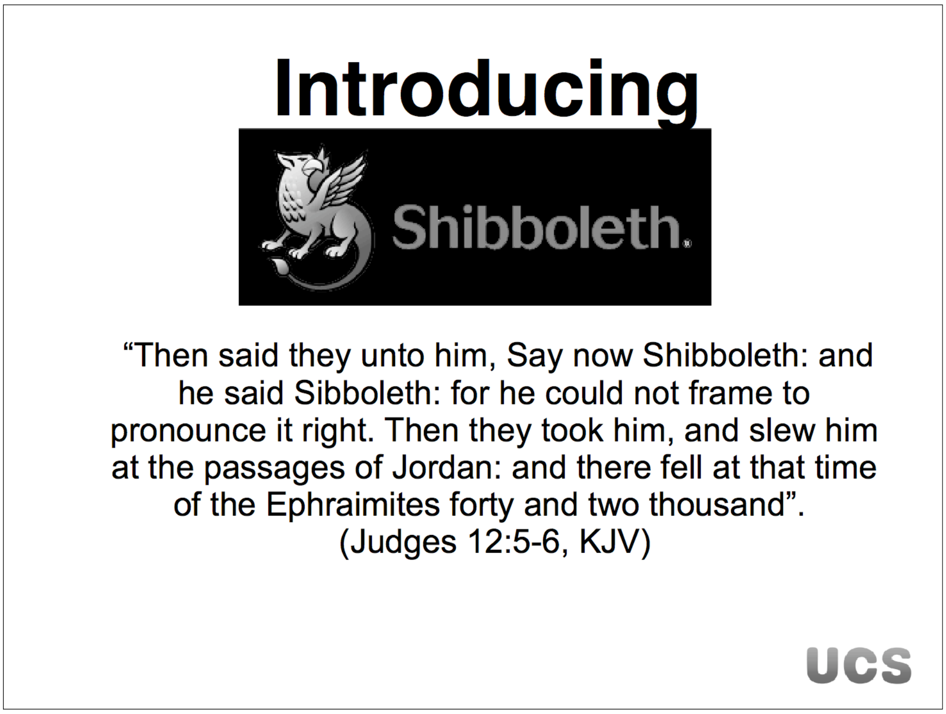

Introducing

Shibboleth

“Then said they unto him, Say now Shibboleth: and he said Sibboleth: for he could not frame to pronounce it right. Then they took him, and slew him at the passages of Jordan: and there fell at that time of the Ephraimites forty and two thousand”.
(Judges 12:5-6, KJV)

UCS

This talk aims to introduce the Shibboleth web authentication/authorization framework and its intended deployment in the UK academic community and the University.

Shibboleth named after an event in the bible where the inability to pronounce 'sh' (as in 'shoe') was used to distinguish members of one tribe from another, with unfortunate consequences for those who couldn't.

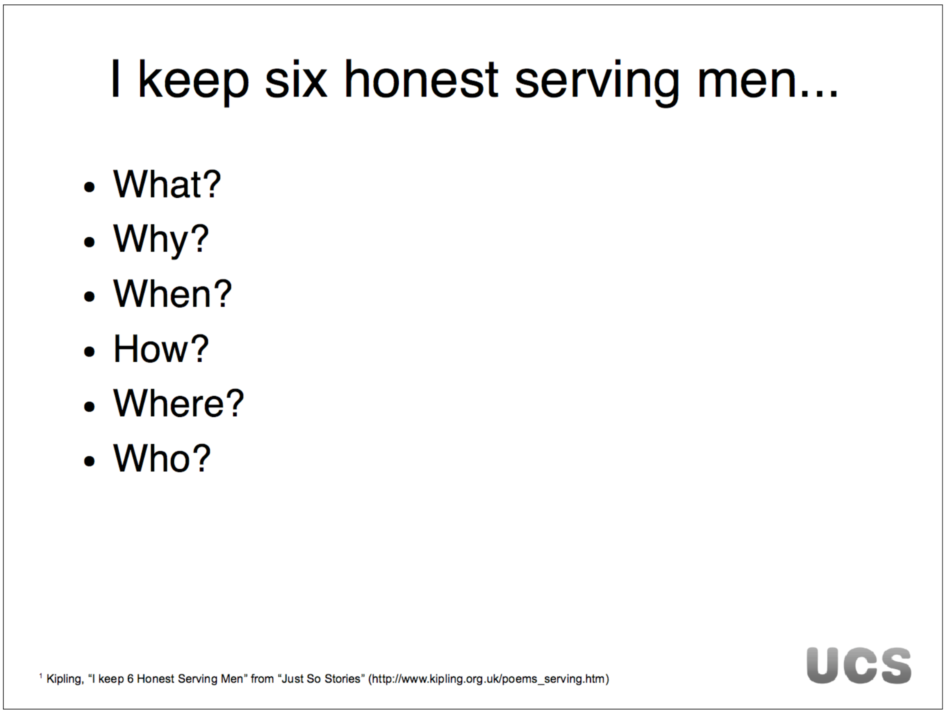

I keep six honest serving men...

- What?
- Why?
- When?
- How?
- Where?
- Who?

[1] Kipling, "I keep 6 Honest Serving Men" from "Just So Stories" (http://www.kipling.org.uk/poems_serving.htm)

Here's a reasonable set of sub-headings for any talk (thanks to Kipling) ...

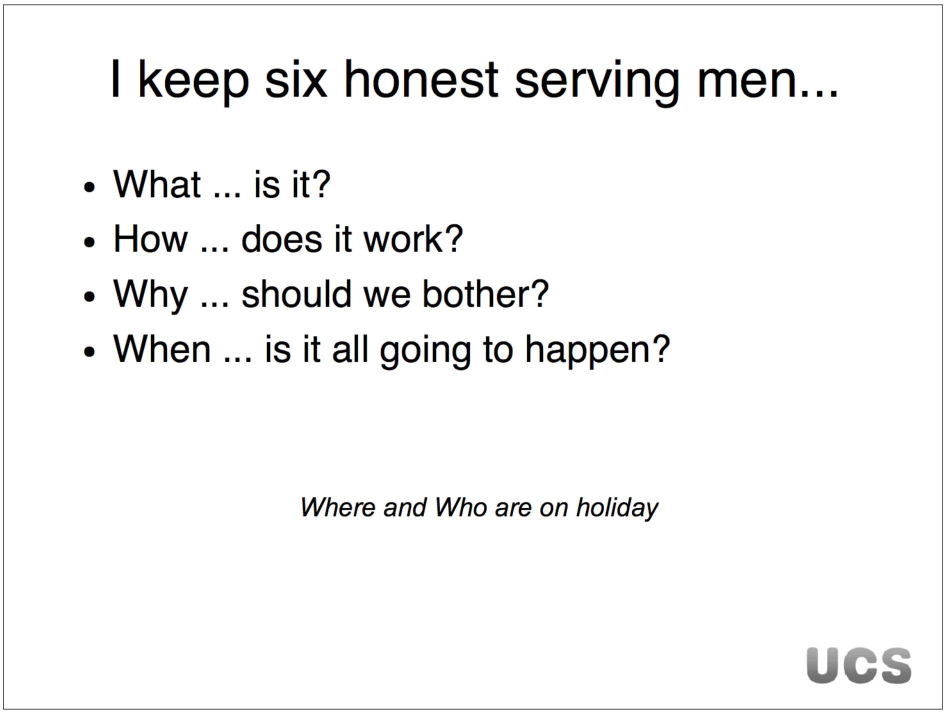

I keep six honest serving men...

- What ... is it?
- How ... does it work?
- Why ... should we bother?
- When ... is it all going to happen?

Where and Who are on holiday

- ... except that I plan to cover them in a different order and ignore 'Where' and 'Who'

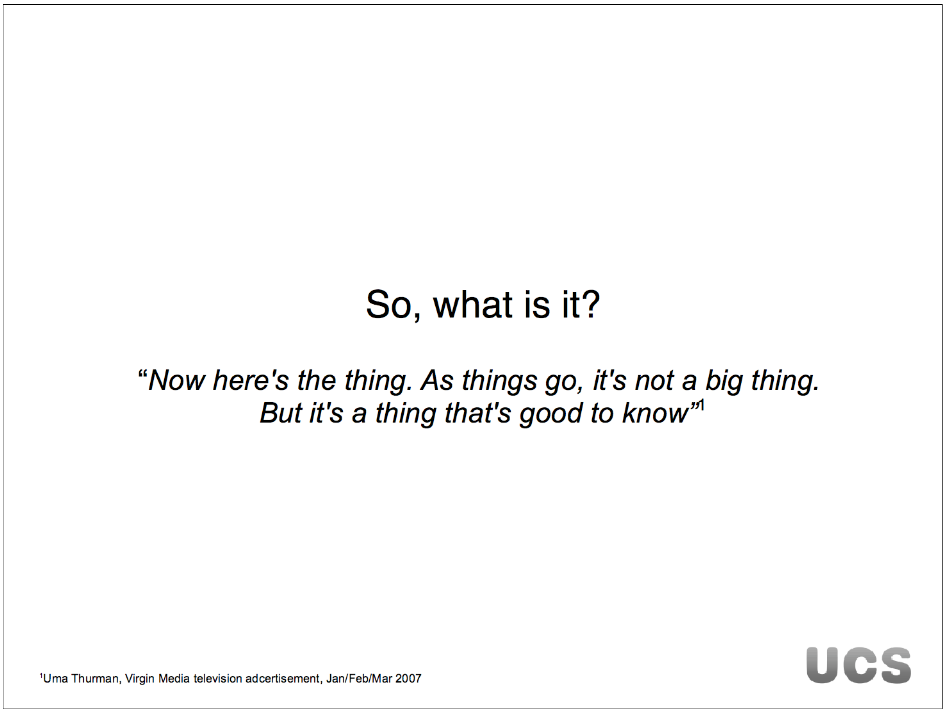

So, what is it?

“Now here's the thing. As things go, it's not a big thing. But it's a thing that's good to know”[1]

[1] Uma Thurman, Virgin Media television adcertisement, Jan/Feb/Mar 2007

“Not a big thing”: while it's perhaps interesting, and something some people will have to understand, Shibboleth is a technology that many of its users will never identify as such. Indeed, we may not call the service 'Shibboleth' at all.

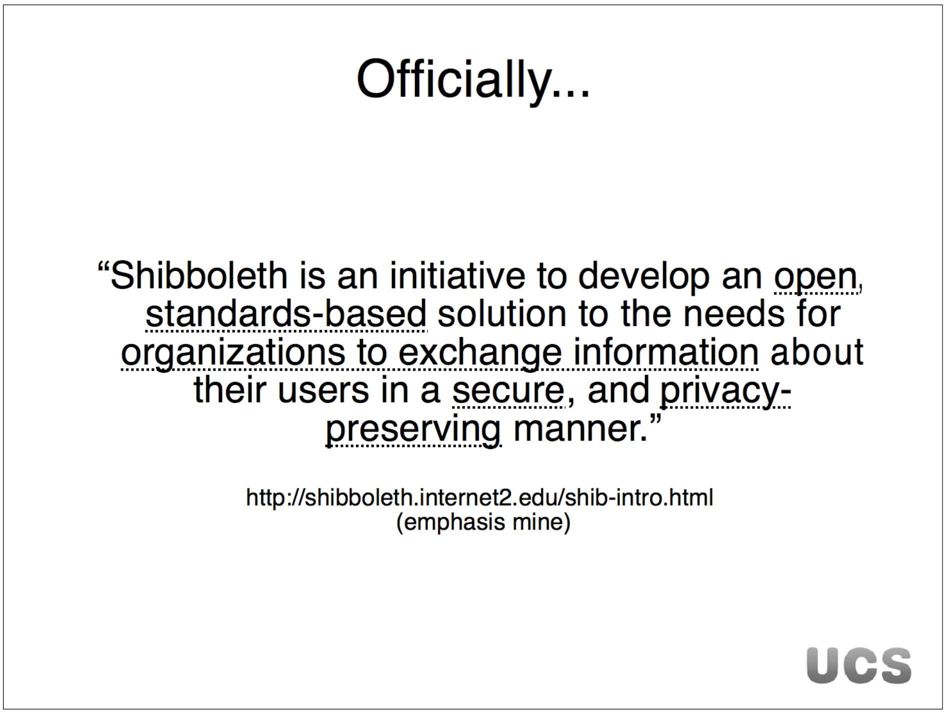

Officially...

“Shibboleth is an initiative to develop an open, standards-based solution to the needs for organizations to exchange information about their users in a secure, and privacy-preserving manner.”

http://shibboleth.internet2.edu/shib-intro.html
(emphasis mine)

Shibboleth was (and is) developed by Internet2 in the US. This is one of the definitions of the Shibboleth project taken from their web site.

Note the underlined bits:

- open
- standards-based
- for organisations to exchange information
- secure
- privacy-preserving

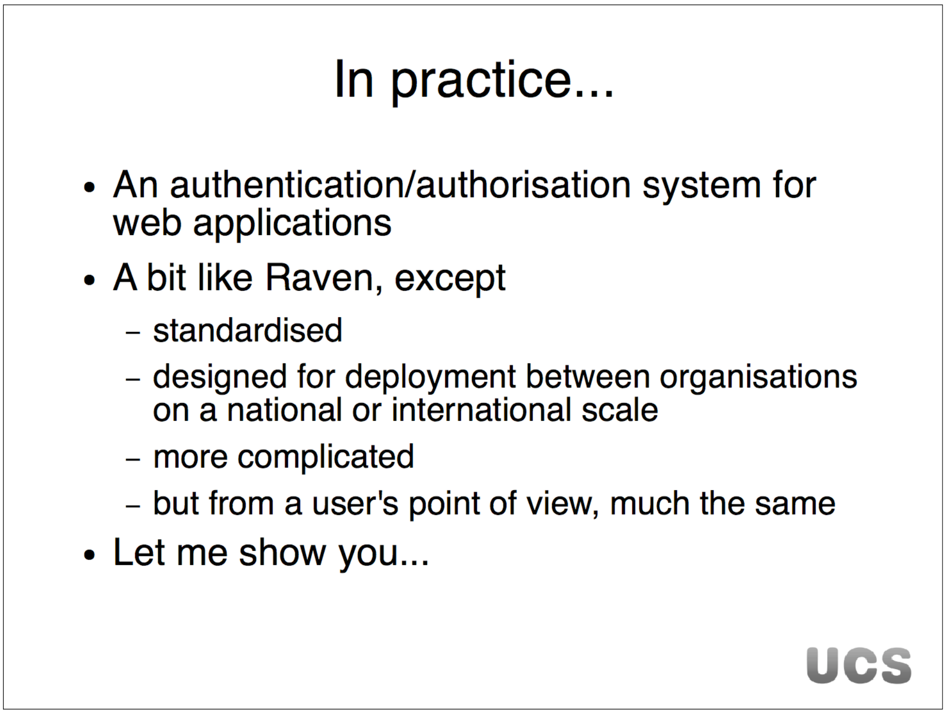

In practice...

- An authentication/authorisation system for web applications
- A bit like Raven, except
  - standardised
  - designed for deployment between organisations on a national or international scale
  - more complicated
  - but from a user's point of view, much the same
- Let me show you...

Shibboleth supports authenticating (i.e. identifying) users of web applications, and making authorisation decisions (i.e. working out what the identified users are allowed to do).

Note that Shibboleth, like Raven, only works in a web environment with a real user driving the browser.

Shibboleth is a standard, which increases the likelihood of people adopting it, and is designed to work across groups of independent organisations (unlike Raven, which only really works within a single organisation).

The additional power of Shibboleth comes at the price of making it all much more complex under the covers.

However many users will hardly notice ...

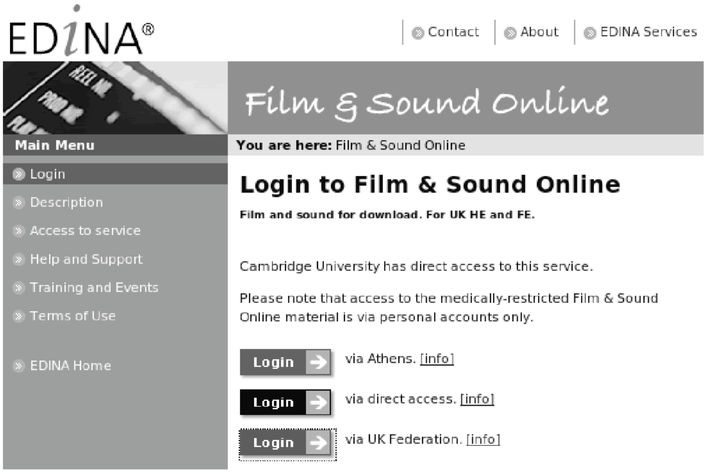

Login to Film & Sound Online

Film and sound for download. For UK HE and FE.

Cambridge University has direct access to this service.

Please note that access to the medically-restricted Film & Sound Online material is via personal accounts only.

via Athens. [info]

via direct access. [info]

via UK Federation. [info]

UCS

This is the login pge for a real electronic resource licensed by the University: Film and Sound Online. This is the login page you get to if you select this resource from the UL's resources page.

From the standard login page, choose 'Login via the UK Federation' to use Shibboleth (more on the Federation later)

Note that the terminology for how you start a Shibboleth login hasn't stabilised yet, and this may be an obstacle to adoption.

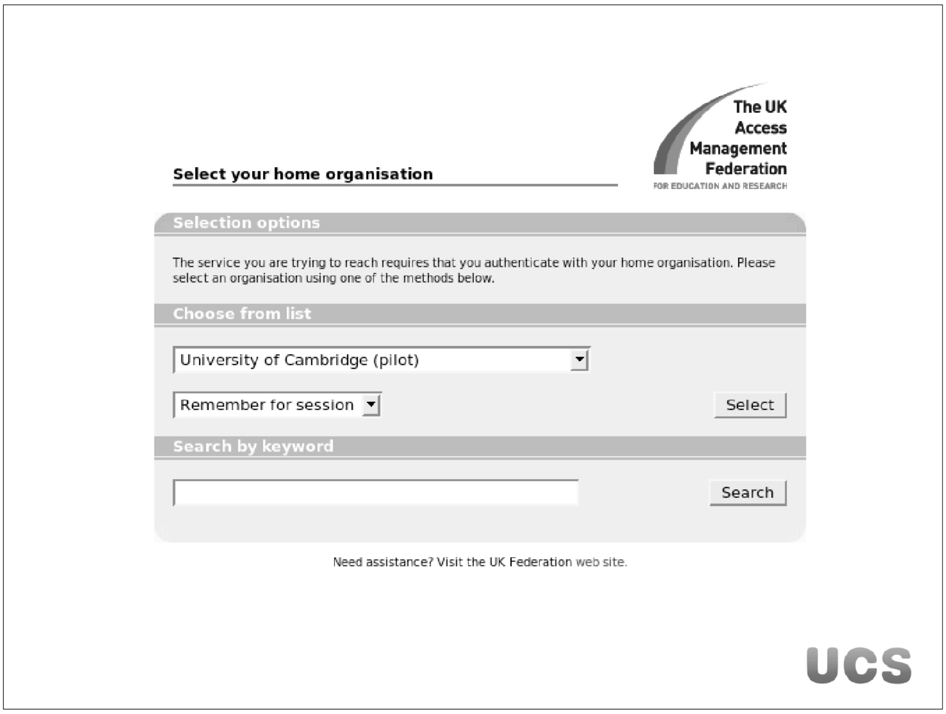

Select your home organisation

Selection options

The service you are trying to reach requires that you authenticate with your home organisation. Please select an organisation using one of the methods below.

Choose from list

University of Cambridge (pilot)

Remember for session

Search by keyword

Need assistance? Visit the UK Federation web site.

The UK Access Management Federation
FOR EDUCATION AND RESEARCH
Select
Search
UCS

In Shibboleth users are always identified by their *home organisation* so that has to be identified somehow.

One way to do this to display a list of possibilities and ask the user to identify one. This is commonly called a 'Where are you from' service, or WAYF.

There are other, perhaps better, ways to address this requirement.

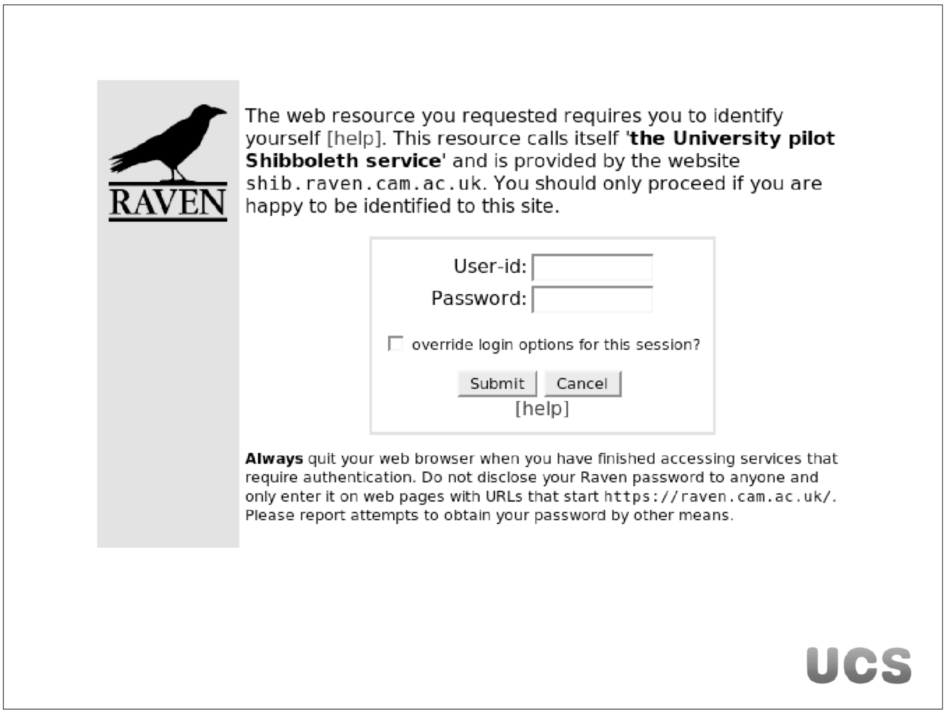

The web resource you requested requires you to identify yourself [help]. This resource calls itself '**the University pilot Shibboleth service**' and is provided by the website `shib.raven.cam.ac.uk`. You should only proceed if you are happy to be identified to this site.

User-id:

Password:

☐ override login options for this session?

Always quit your web browser when you have finished accessing services that require authentication. Do not disclose your Raven password to anyone and only enter it on web pages with URLs that start `https://raven.cam.ac.uk/`. Please report attempts to obtain your password by other means.

The next thing the user sees is their home organisation's web login service. For University of Cambridge users, this is Raven

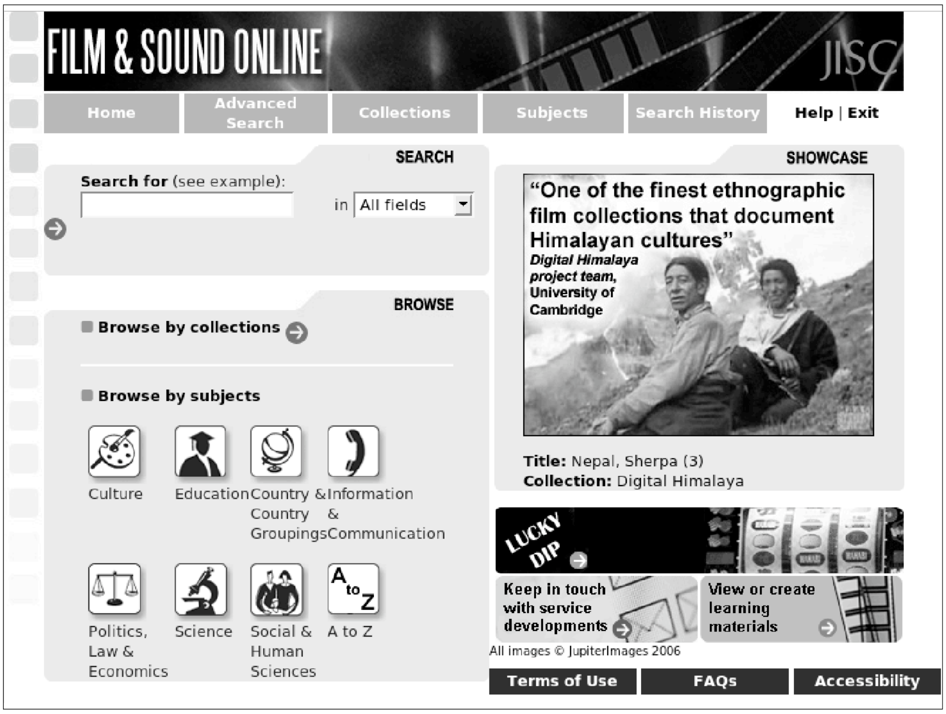

"One of the finest ethnographic film collections that document Himalayan cultures"
Digital Himalaya project team, **University of Cambridge**

Title: Nepal, Sherpa (3)
Collection: Digital Himalaya

... and the next thing the user sees is the resource they wanted.

The whole sequence is much the same as you'd see if you accessed a Raven-protected site inside the University.

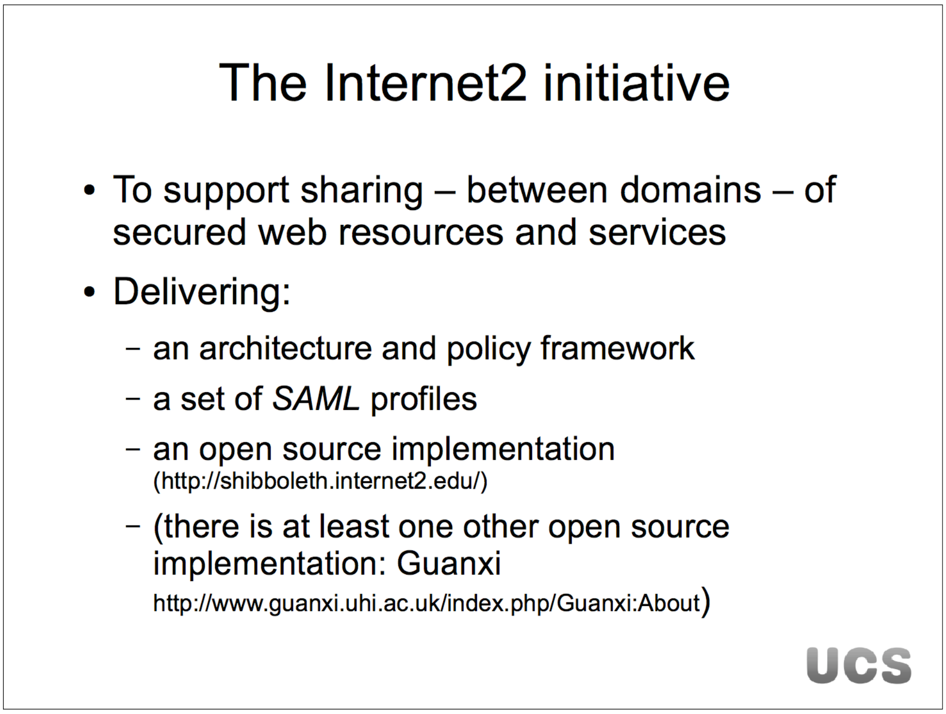

The Internet2 initiative

- To support sharing – between domains – of secured web resources and services
- Delivering:
  - an architecture and policy framework
  - a set of *SAML* profiles
  - an open source implementation (http://shibboleth.internet2.edu/)
  - (there is at least one other open source implementation: Guanxi http://www.guanxi.uhi.ac.uk/index.php/Guanxi:About)

Internet2's Shibboleth project aims to support inter-domain web resource sharing. It actually represents three separate developments:

- a set of architectures and policies within which this sharing can take place
- a set of profiles of the *Security Assertion Markup Language* to support this
- an example open-source implementation

There is at least one other Shibboleth implementation: Guanxi from the UK's University of the Highlands and Islands

Shibboleth relies on the Security Assertion Markup Language (SAML), an XML dialect for exchanging authentication and authorisation data [*authentication* - the process of proving who you are; *authorisation* - the process of working out if you can do what you are trying to do]

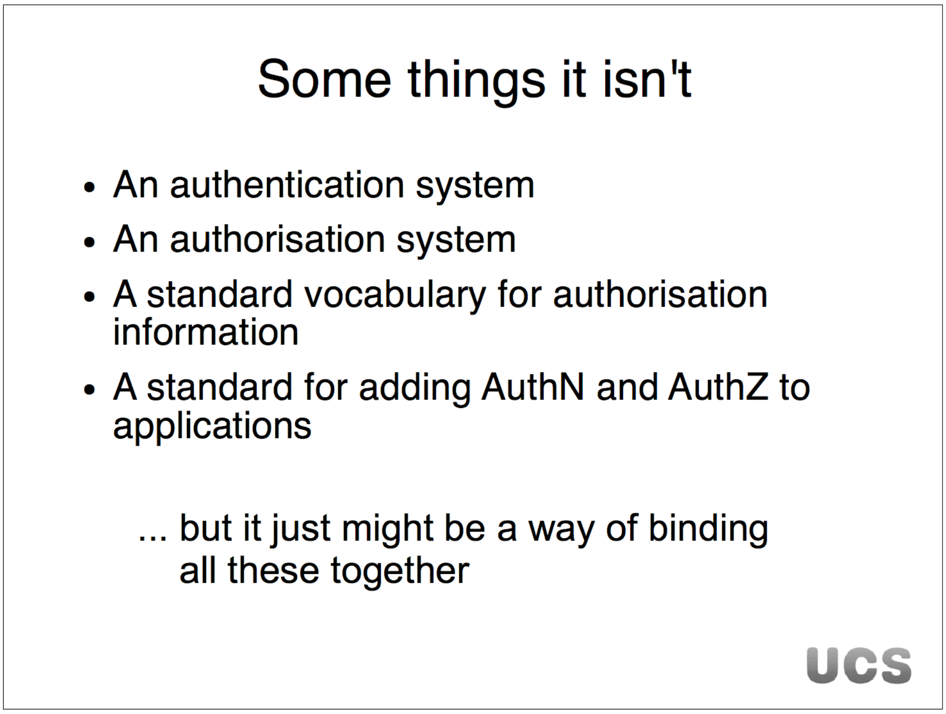

Some things it isn't

- An authentication system
- An authorisation system
- A standard vocabulary for authorisation information
- A standard for adding AuthN and AuthZ to applications

... but it just might be a way of binding all these together

UCS

Despite it's capabilities, there are a surprising number of things that Shibboleth itself doesn't specify:

- It doesn't specify how users are actually authenticated – in practice existing systems, such as Raven, are normally reused
- It doesn't say how authorisation decisions should be taken or implemented
- It doesn't even define how information needed for authorisation should be represented
- It doesn't provide any guidance on how its authentication or authorisation services should or could be added to web applications

BUT it does provide a framework for binding together existing implementations of all these into something that proves to be useful.

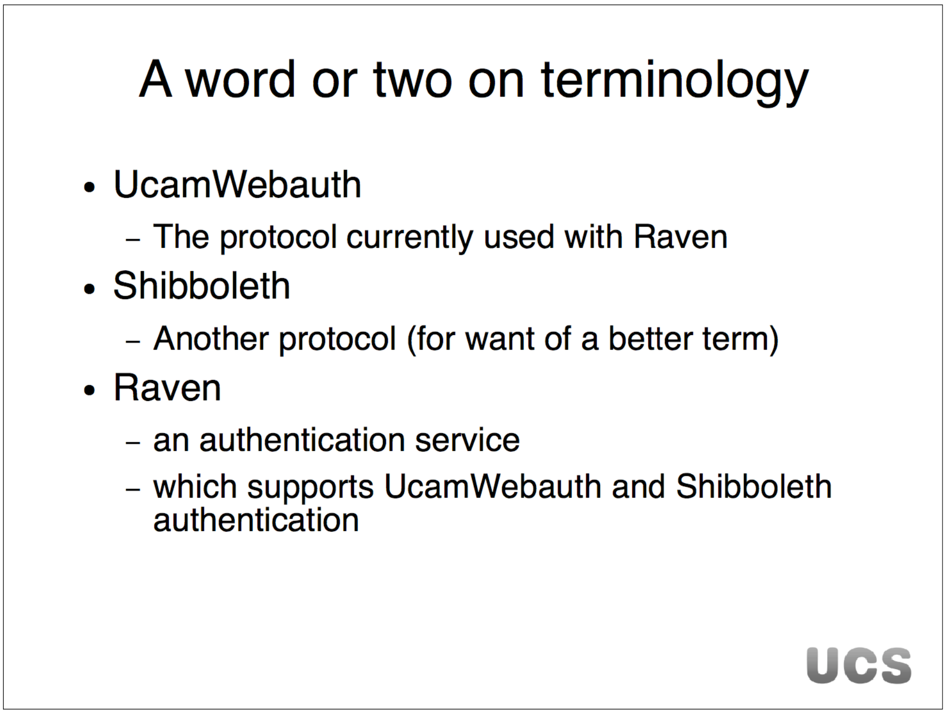

A word or two on terminology

- UcamWebauth
  - The protocol currently used with Raven
- Shibboleth
  - Another protocol (for want of a better term)
- Raven
  - an authentication service
  - which supports UcamWebauth and Shibboleth authentication

Everyone (me included) have got used to saying 'Raven' when we mean any of the separate components that make up the Raven service (as in 'I logged in with Raven')

Adding Shibboleth to the Raven service will force many of us to be more careful about the terms we use.

How does it work?

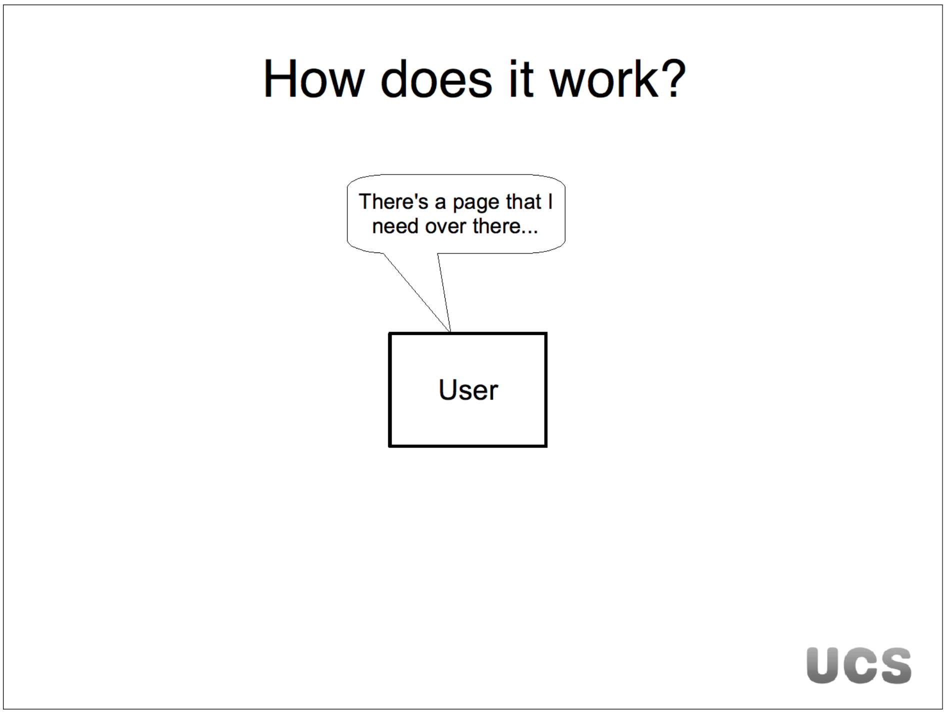

How does it work?

We start off with a user who has identified a link to some information that they think will be useful to them. This link could be on a library website, on a handout, in a VLE, on the back of a bus, etc.

There is nothing Shibboleth-special about this link – it could just as well be used by someone who was going to authenticate using some non-Shibboleth system.

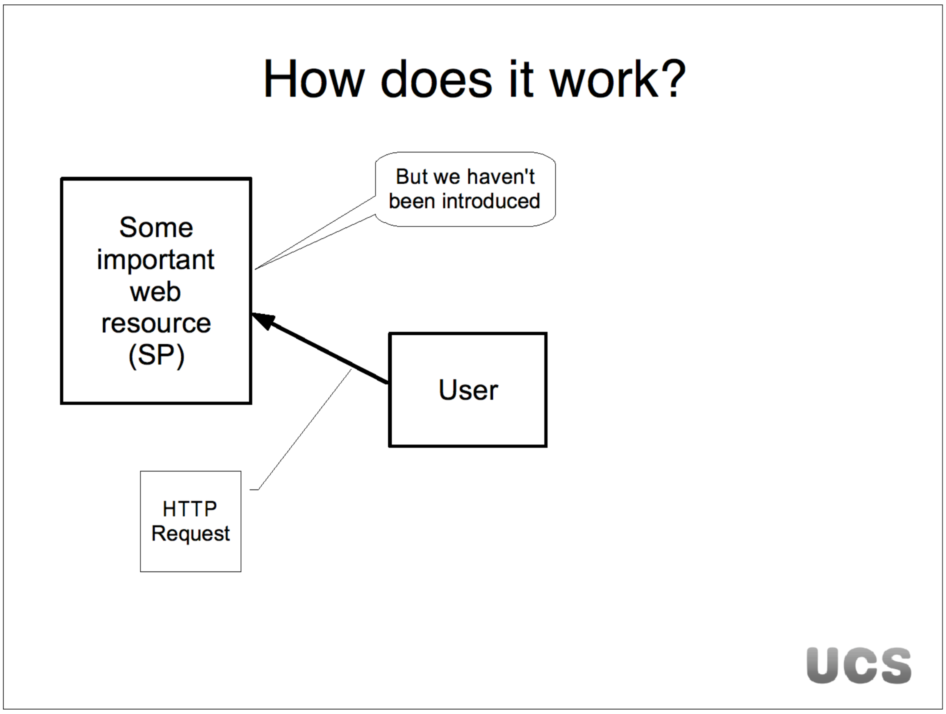

How does it work?

Some important web resource (SP)

But we haven't been introduced

User

HTTP Request

The user's browser follows the link, contacts the site, and requests the resource.

Web resources protected by Shibboleth are called *Service Providers* (SPs).

The SP can't make any authorisation decisions because it knows nothing about this user.

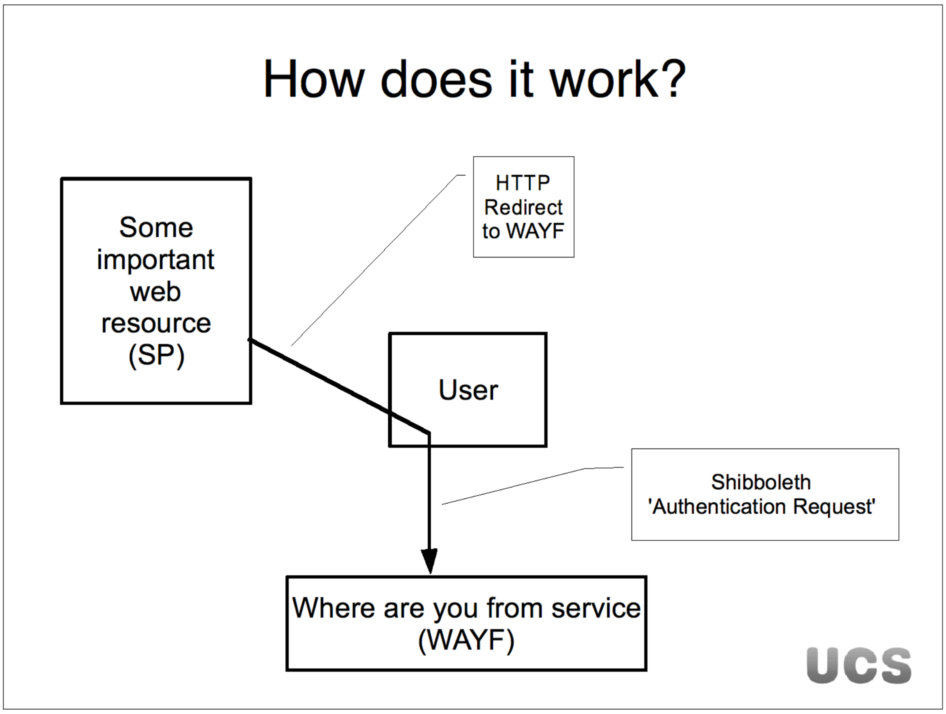

How does it work?

Some important web resource (SP)

HTTP Redirect to WAYF

User

Shibboleth 'Authentication Request'

Where are you from service (WAYF)

The SP constructs a Shibboleth *Authentication Request* message. This includes, amongst other things, the identity of the SP and the URL of the resource the user wants to access.

This needs to be processed by the user's *home organisation*, which in Shibboleth is called an *Identity Provider* (IdP) ...

... but since the SP doesn't know (or really care) where this is, it sends this message, and the user, to a *Where are you from service* (WAYF).

It does this using an 'HTTP Redirect' – essentially telling the user's browser to retrieve a new URL. The Authentication request is encoded in this URL.

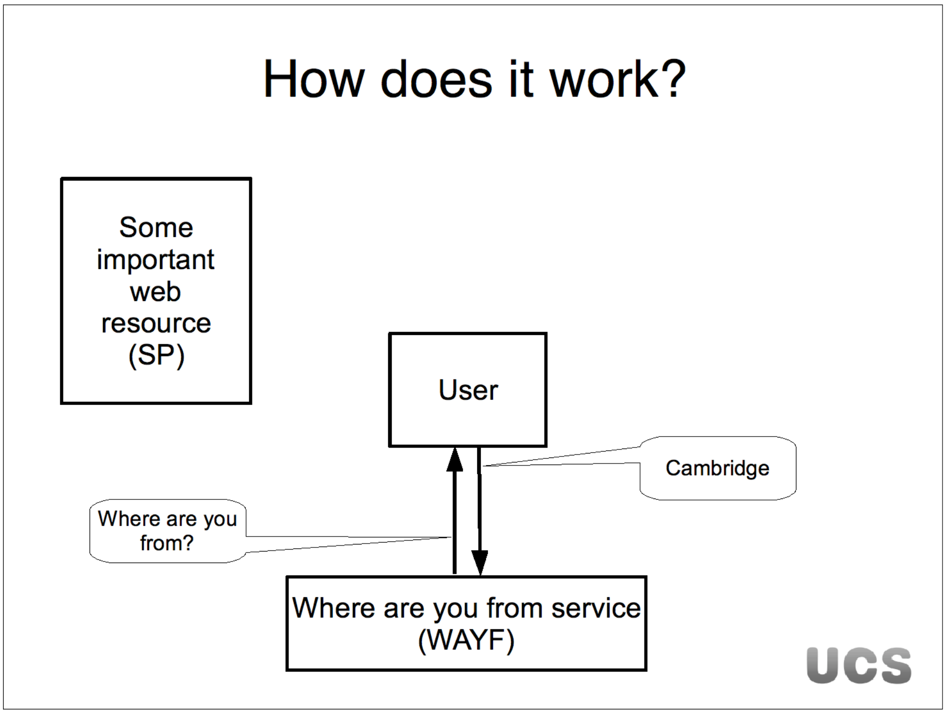

How does it work?

Some important web resource (SP)

User

Cambridge

Where are you from?

Where are you from service (WAYF)

The user and the WAYF interact to identify the user's IdP.

There are lots of ways that this might happen (especially on a second and subsequent visit to the WAYF) but as a backstop the WAYF can ask the user to select from a list of all the IdPs it knows about.

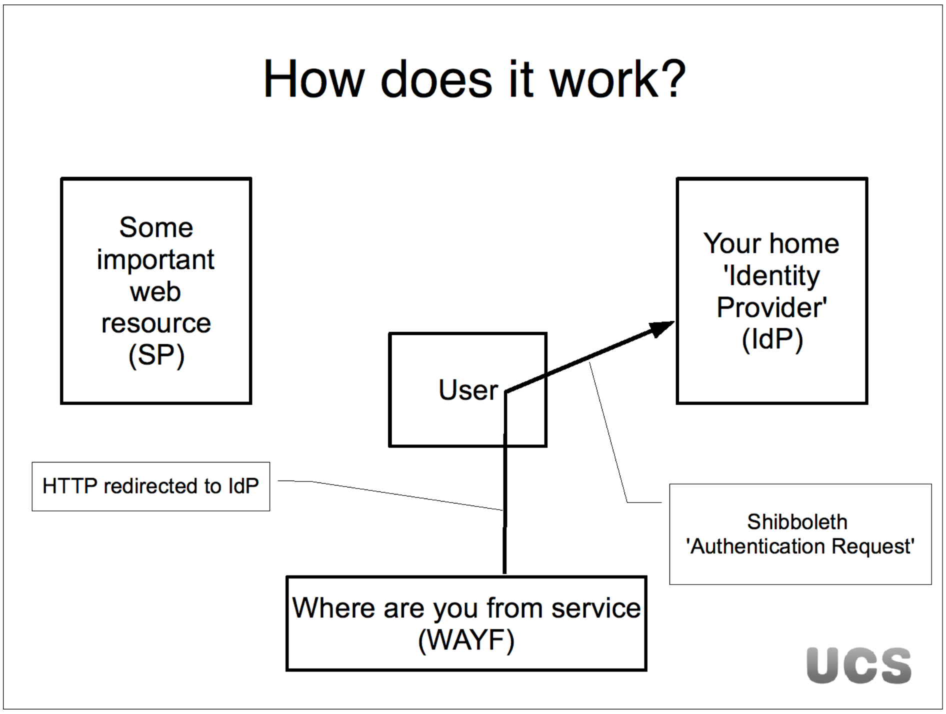

How does it work?

Some important web resource (SP)

User

Your home 'Identity Provider' (IdP)

HTTP redirected to IdP

Shibboleth 'Authentication Request'

Where are you from service (WAYF)

Having identified the user's IdP, the WAYF redirects the user's browser there, complete with the original Authentication request.

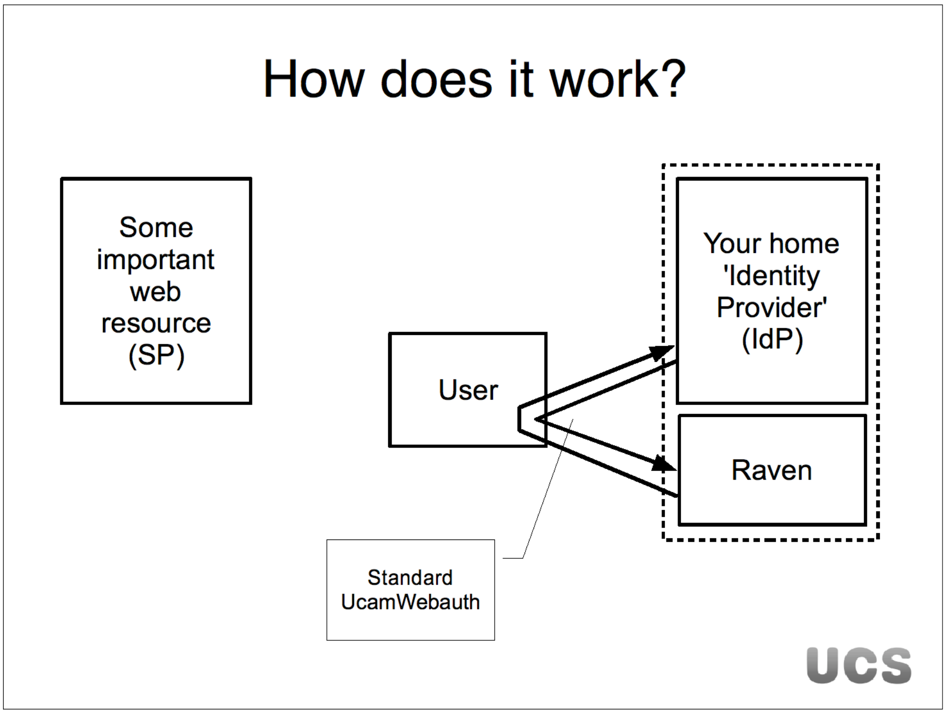

How does it work?

Some important web resource (SP)
User
Your home 'Identity Provider' (IdP)
Raven
Standard UcamWebauth
UCS

The IdP uses some apropriate means to identify the user. This is probably going to involve a campus SSO system - we use Raven (which in turn uses the UcamWebauth protocol 'on the wire')

The outcome of this is that the IdP now knows the identity of the browser user.

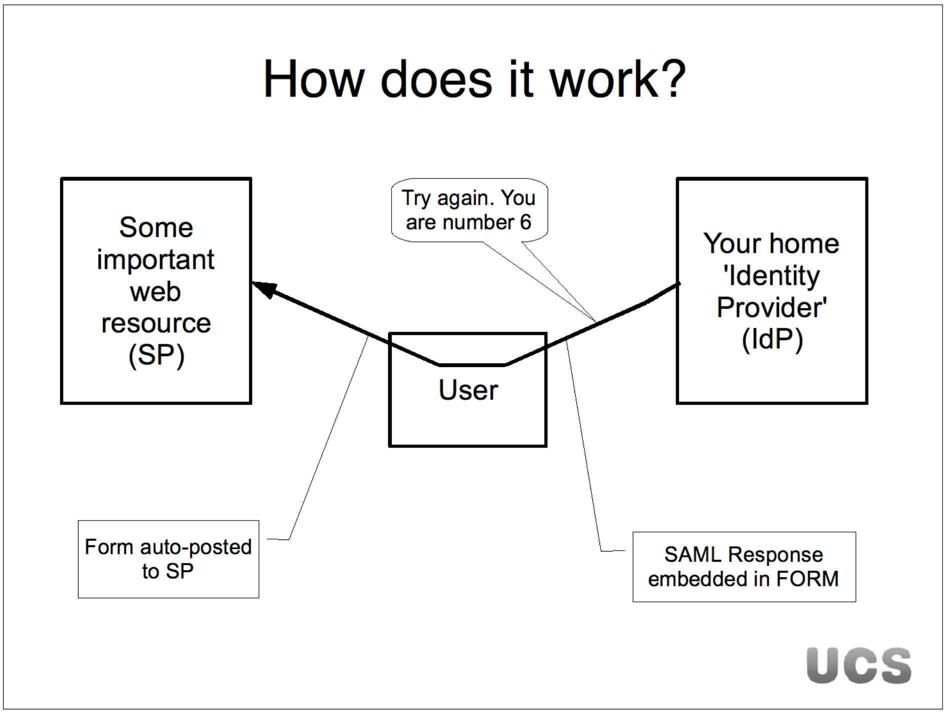

How does it work?

Some important web resource (SP)

Try again. You are number 6

Your home 'Identity Provider' (IdP)

User

Form auto-posted to SP

SAML Response embedded in FORM

Once the IdP knows who the user is, it sends the user's browser back to the SP with a SAML *Authentication Assertion* message.

This message contains no real information about the user but does assert that the IdP has successfully identified them, how it did so, when it happened, etc. The IdP also copies the URL that the user originally requested into the response from the request.

This message can't be sent as a HTTP Redirect because it is too big, so it is sent by embedding it in a form field and then using JavaScript to automatically submit the form (or getting the user to click a button). The form is submitted to the SPs *Assertion Consumer Service* (the location of which was also included in the authentication request).

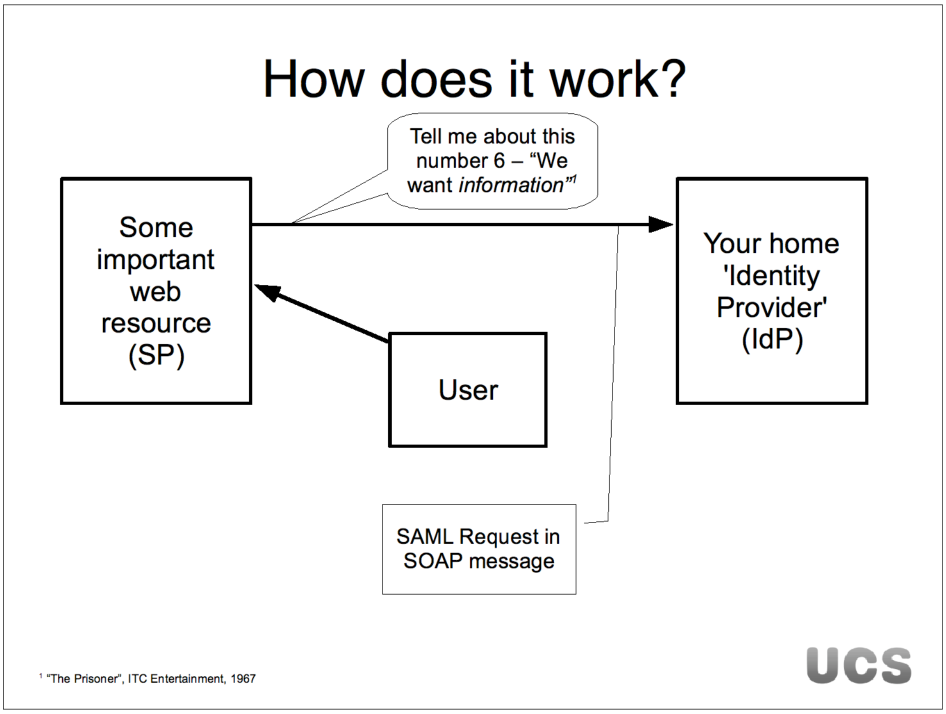

How does it work?

Some important web resource (SP)

Your home 'Identity Provider' (IdP)

User

Tell me about this number 6 – “We want *information*”[1]

SAML Request in SOAP message

[1] “The Prisoner”, ITC Entertainment, 1967

The SP still doesn't know anything about the user, except that they managed to authenticate at this particular IdP.

The SP sends a SAML *Attribute Query* direct to the IdP's *Attribute Authority* service asking for more information. It can, but typically doesn't, say what it wants to know.

The query is carried in a SOAP message over HTTP.

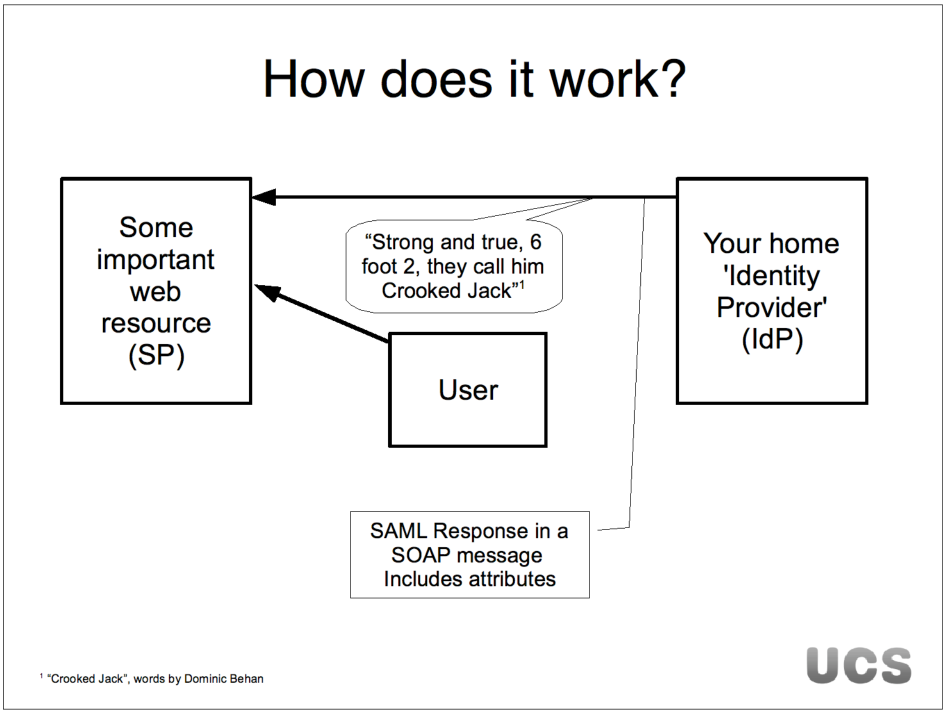

How does it work?

Some important web resource (SP)

"Strong and true, 6 foot 2, they call him Crooked Jack"[1]

Your home 'Identity Provider' (IdP)

User

SAML Response in a SOAP message Includes attributes

[1] "Crooked Jack", words by Dominic Behan

The IdP works out what information it is willing to release, based on what it has available, who the user is, and who the SP is.

It sends this information back to the SP in a SOAP message encapsulating a SAML *Attribute Assertion*.

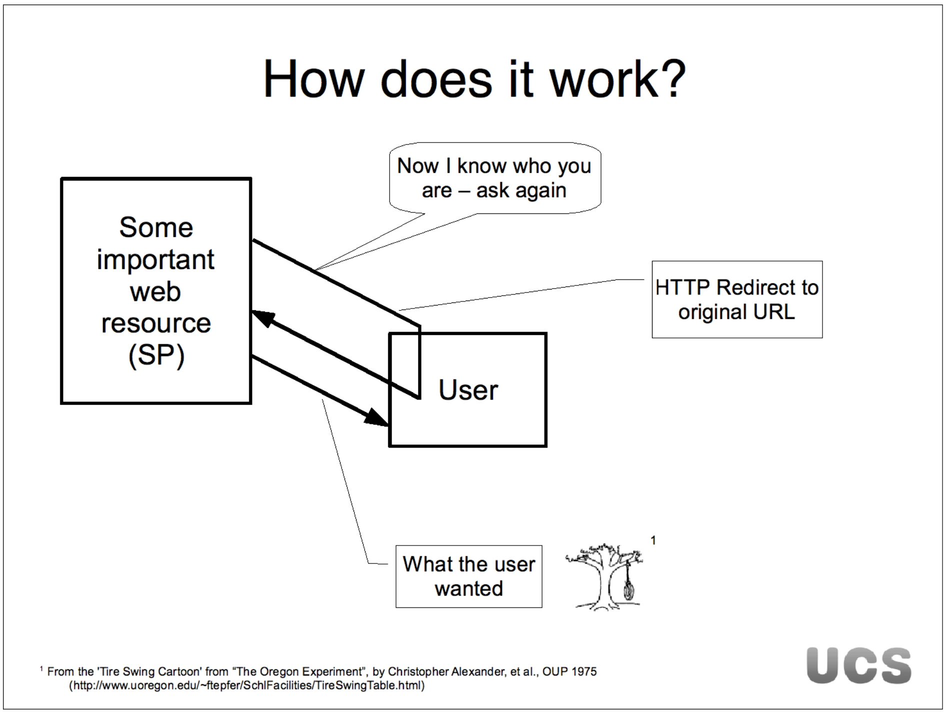

How does it work?

Now I know who you are – ask again
Some important web resource (SP)
HTTP Redirect to original URL
User
What the user wanted

[1] From the 'Tire Swing Cartoon' from "The Oregon Experiment", by Christopher Alexander, et al., OUP 1975 (http://www.uoregon.edu/~ftepfer/SchlFacilities/TireSwingTable.html)

The SP decides what it wants to do, based on the attribute values it receives.

Whatever it decides, it redirects the user to a new page – if everything has gone to plan this will be the resource the user requested in the first place.

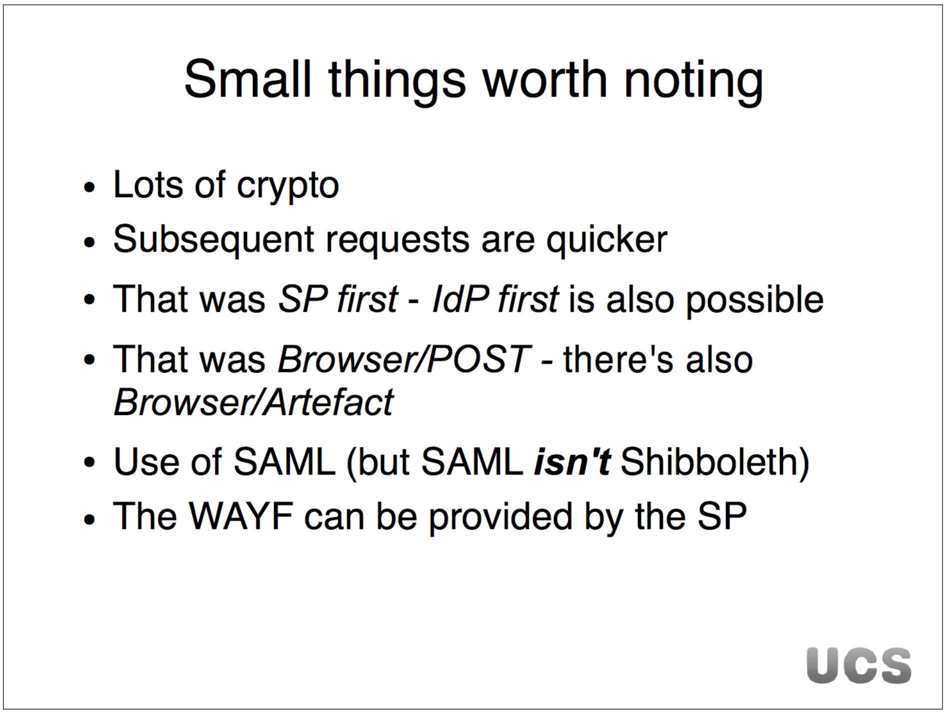

Small things worth noting

- Lots of crypto
- Subsequent requests are quicker
- That was *SP first* - *IdP first* is also possible
- That was *Browser/POST* - there's also *Browser/Artefact*
- Use of SAML (but SAML ***isn't*** Shibboleth)
- The WAYF can be provided by the SP

The authentication assertion is signed to prevent tampering. The attribute query and attribute assertion travel over an SSL/TLS connection which allows both ends to be strongly authenticated.

The whole exchange can be quicker for subsequent authentications – WAYFs typically provide short cuts for accessing recent IdPs, Raven only requires one login per session, etc.

Other message sequences are possible to achieve the same thing – we described an *SP-first* flow, but it is also possible to have an *IdP-first* flow where the user visits the IdP first, authenticate, and only then contact the SP. Another possibility is to avoid the use of forms and JavaScript (the SAML *Browser/POST profile*) by redirecting the user back to the SP with a small random token called an artefact which the SP then uses to collect the full authentication assertion over SOAP. This is know as the SAML *Browser/Artefact profile*.

Note the reliance on SAML, but remember that Shibboleth is just one of many possible applications that use SAML.

The WAYF can be provided by the SP – this can result in a cleaner interface since the SP knows who its customers are (e.g. Science Direct)

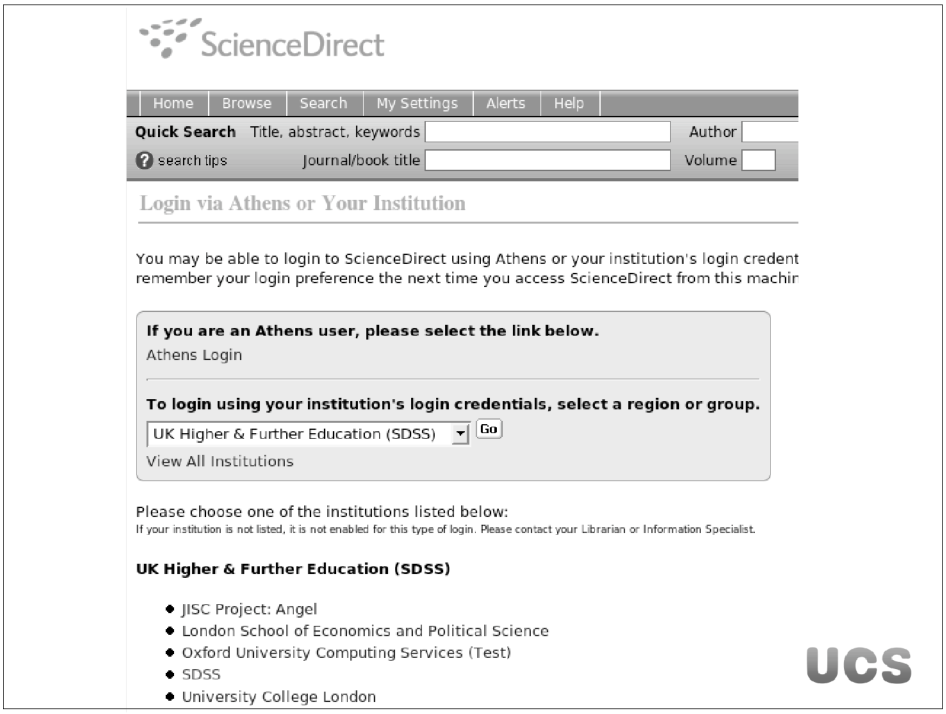

Login via Athens or Your Institution

You may be able to login to ScienceDirect using Athens or your institution's login credent
remember your login preference the next time you access ScienceDirect from this machir

If you are an Athens user, please select the link below.
Athens Login

To login using your institution's login credentials, select a region or group.
UK Higher & Further Education (SDSS)
View All Institutions

Please choose one of the institutions listed below:
If your institution is not listed, it is not enabled for this type of login. Please contact your Librarian or Information Specialist.

UK Higher & Further Education (SDSS)

- JISC Project: Angel
- London School of Economics and Political Science
- Oxford University Computing Services (Test)
- SDSS
- University College London

Here is Science Direct providing its own WAYF service, rather than using a default provided by a Federation – it only lists institutions that have access.

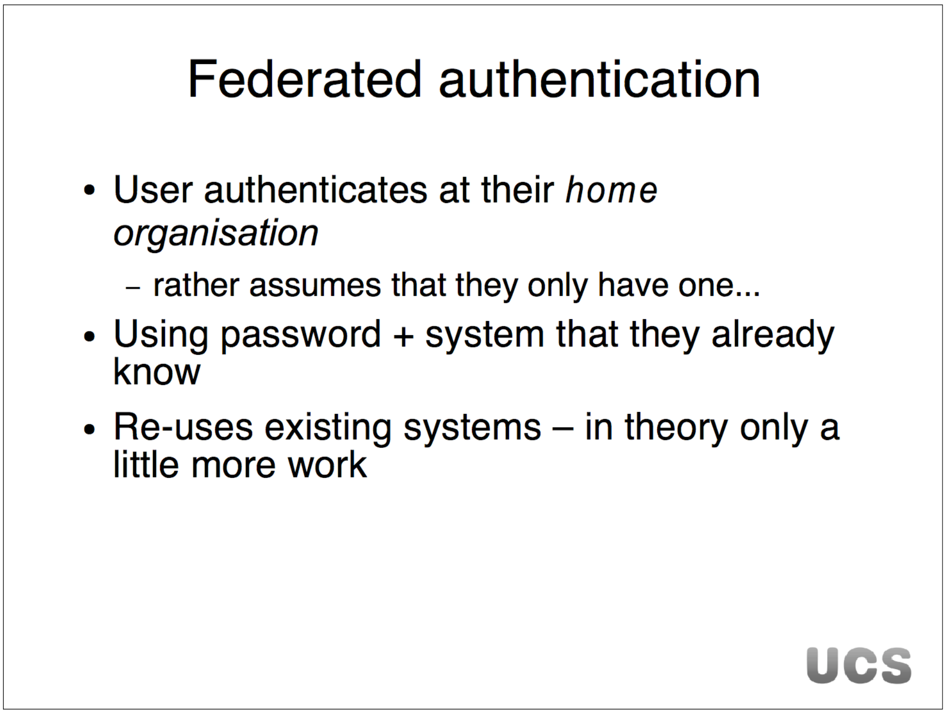

Federated authentication

- User authenticates at their *home organisation*
  - rather assumes that they only have one...
- Using password + system that they already know
- Re-uses existing systems – in theory only a little more work

Shibboleth implements *Federated Authentication*, in which one organisation chooses to trust an authentication decision made by someone else.

This sort of authentication should make life easier for users since it uses something that they (should) already know, rather than a new user name and password. It also saves the SP from having to issue and manage its own authentication credentials.

Shibboleth (currently) rather assumes people only have a single *home organisation* – you can use more than one, but only one at a time. This may eventually be problem for people with multiple affiliations who want to take advantage of all their entitlements simultaneously.

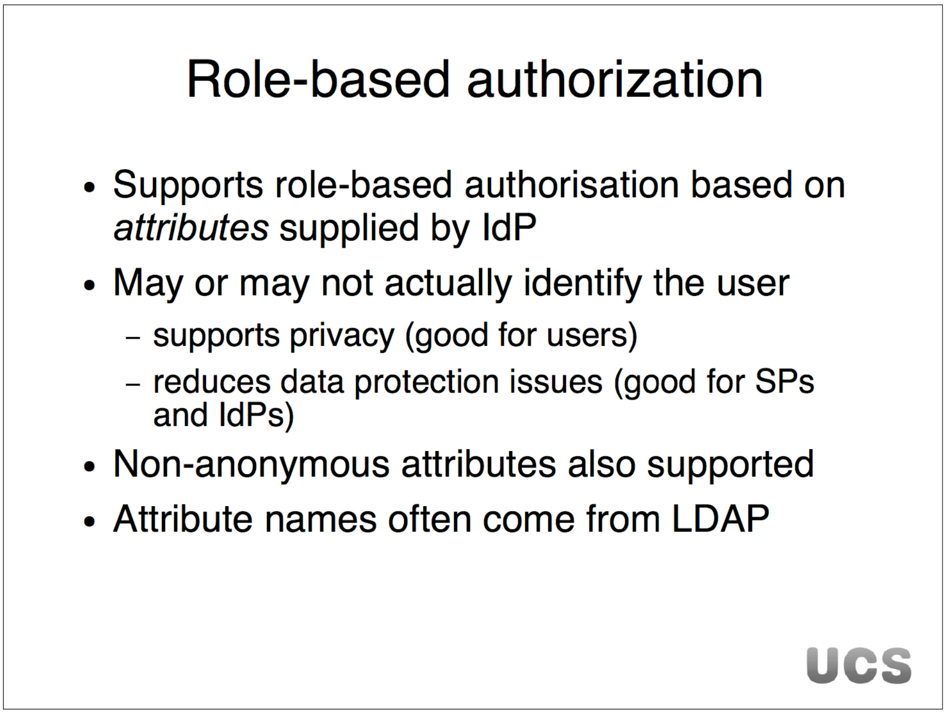

Role-based authorization

- Supports role-based authorisation based on *attributes* supplied by IdP
- May or may not actually identify the user
  - supports privacy (good for users)
  - reduces data protection issues (good for SPs and IdPs)
- Non-anonymous attributes also supported
- Attribute names often come from LDAP

Shibboleth supports *role-based authorisation*, in which what you can do depends not directly on who you are but on one or more bits of information about you. These bits of information are called *Attributes* and are supplied by your IdP*.*

These attributes don't always need to include your real-world identity. For example to access a resource site-licensed by the University of Cambridge it may be sufficient just to show that you are a member of the University. This is good for you since it protects your privacy, and is good for IdPs and SPs since it saves them from having to process, and therefore protect, *personal data*.

An IdP can maintain sufficient information to associate each authentication assertion with the user on who's behalf it was made. This allows potential abuse to be investigated even though the SP doesn't know who might have done it.

Where required, Shibboleth can also assert non anonymous attributes – e.g. name and email address to support account provisioning at the SP.

While Shibboleth oftern uses names and meanings from LDAP schema, you don't need an LDAP directory to use Shibboleth [LDAP: Lightweight Directory Access Protocol]

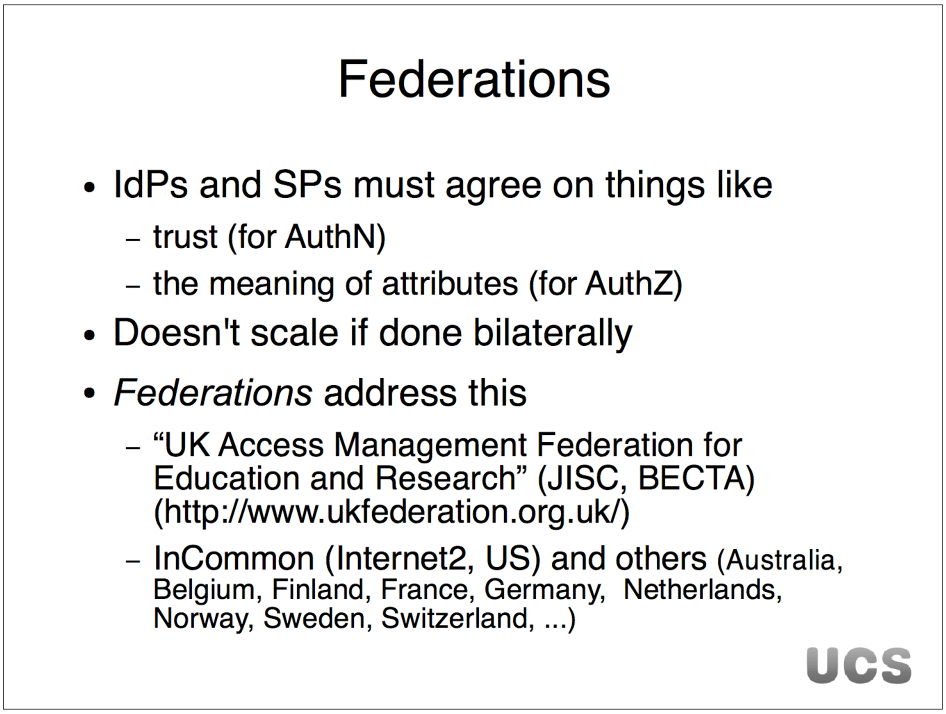

Federations

- IdPs and SPs must agree on things like
  - trust (for AuthN)
  - the meaning of attributes (for AuthZ)
- Doesn't scale if done bilaterally
- *Federations* address this
  - “UK Access Management Federation for Education and Research” (JISC, BECTA) (http://www.ukfederation.org.uk/)
  - InCommon (Internet2, US) and others (Australia, Belgium, Finland, France, Germany, Netherlands, Norway, Sweden, Switzerland, ...)

For federated authentication and authorisation to work, IdPs and SPs have to agree about many things: trust, attribute vocabularies, locations of services, key material, etc. Federations are a largely administrative arrangement to save every SP from having to negotiate agreements with every IdP they want to use.

Federations typically have rules about what particular data items mean, how reliable IdP data needs to be, what SPs should (and shouldn't) do with it, etc. To join the federation, members must agree to abide by the rules.

The *UK Access Management Federation for Education and Research*, commonly '*the UK Federation*', has been established jointly by the Joint Information Systems Committee (JISC) on behalf of FE, and BECTA on behalf of schools. Similar federations are being established in the US, Europe, and elsewhere.

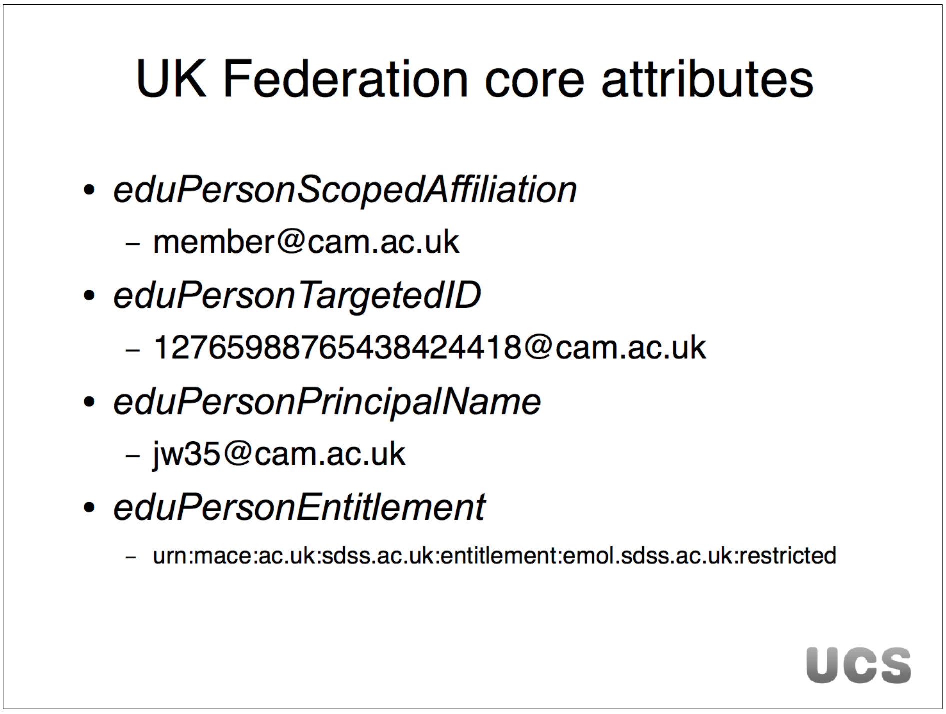

UK Federation core attributes

- *eduPersonScopedAffiliation*
  - member@cam.ac.uk
- *eduPersonTargetedID*
  - 12765988765438424418@cam.ac.uk
- *eduPersonPrincipalName*
  - jw35@cam.ac.uk
- *eduPersonEntitlement*
  - urn:mace:ac.uk:sdss.ac.uk:entitlement:emol.sdss.ac.uk:restricted

UCS

One of the functions of a federation is to define attribute vocabularies. The UK Federation defines 4 *core attributes* that IdPs should aim to provide. SPs are encouraged not to require anything else though they are free to use additional attributes if they are available.

Three of these core attributes are essentially anonymous and so don't represent *personal data*. However one, eduPersonPrincipleName, isn't and does - the federation recommends that it should only be released if it's really required and only when appropriate legal protection is in place.

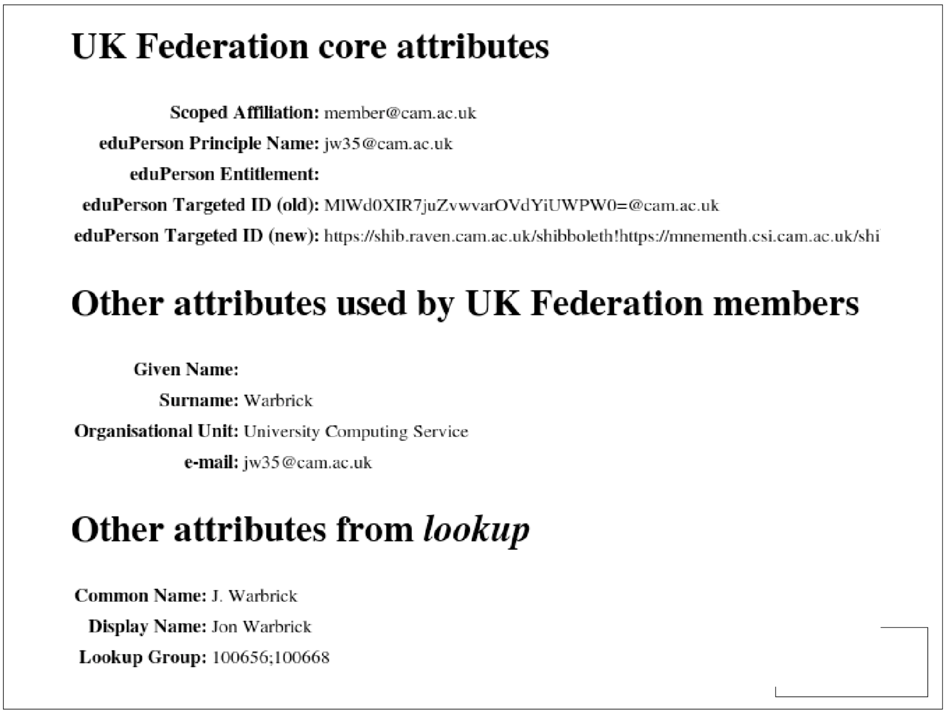

UK Federation core attributes

Scoped Affiliation: member@cam.ac.uk
eduPerson Principle Name: jw35@cam.ac.uk
eduPerson Entitlement:
eduPerson Targeted ID (old): MlWd0XIR7juZvwvarOVdYiUWPW0=@cam.ac.uk
eduPerson Targeted ID (new): https://shib.raven.cam.ac.uk/shibboleth!https://mnementh.csi.cam.ac.uk/shi

Other attributes used by UK Federation members

Given Name:
Surname: Warbrick
Organisational Unit: University Computing Service
e-mail: jw35@cam.ac.uk

Other attributes from *lookup*

Common Name: J. Warbrick
Display Name: Jon Warbrick
Lookup Group: 100656;100668

This is a display of some attribute values being supplied to a SP within the University – many of these would not normally be disclosed outside the University.

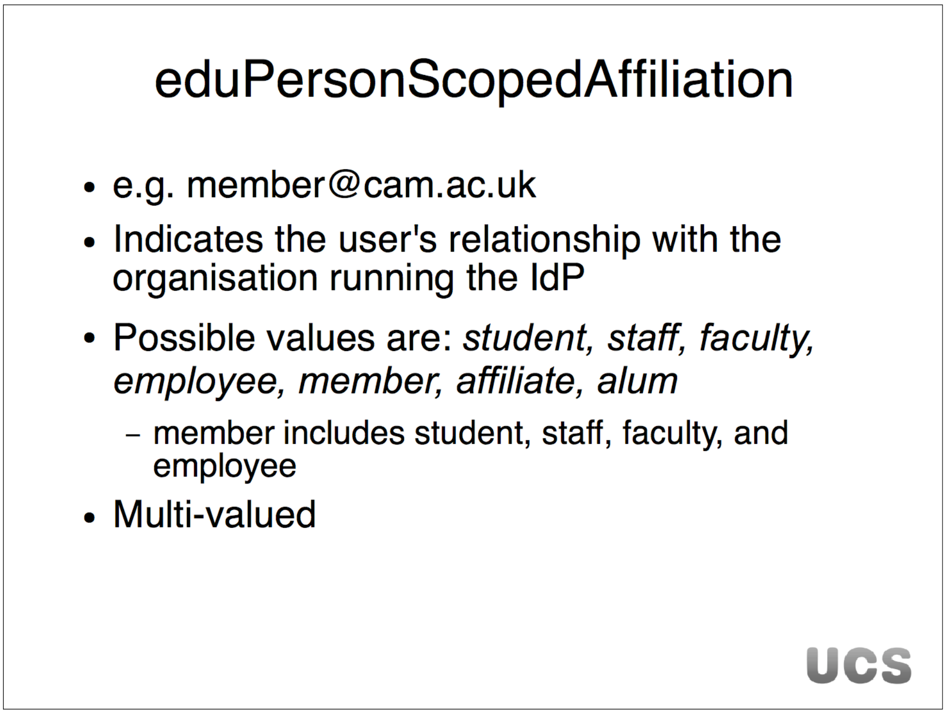

eduPersonScopedAffiliation

- e.g. member@cam.ac.uk
- Indicates the user's relationship with the organisation running the IdP
- Possible values are: *student, staff, faculty, employee, member, affiliate, alum*
  - member includes student, staff, faculty, and employee
- Multi-valued

eduPersonScopedAffiliation (ePSA) is a potentially multi-valued attribute reflecting the authenticated user's association with the IdP's organisation. The possible data values are fixed, and are *student, staff, faculty, employee, member, affiliate*, and *alum*. It is expected that *member* will include anyone who is also *student, staff, faculty*, or *employee.*

It is expected that for many applications, examination of this attribute should be sufficient to determine whether the user has sufficient privilege to access the resource.

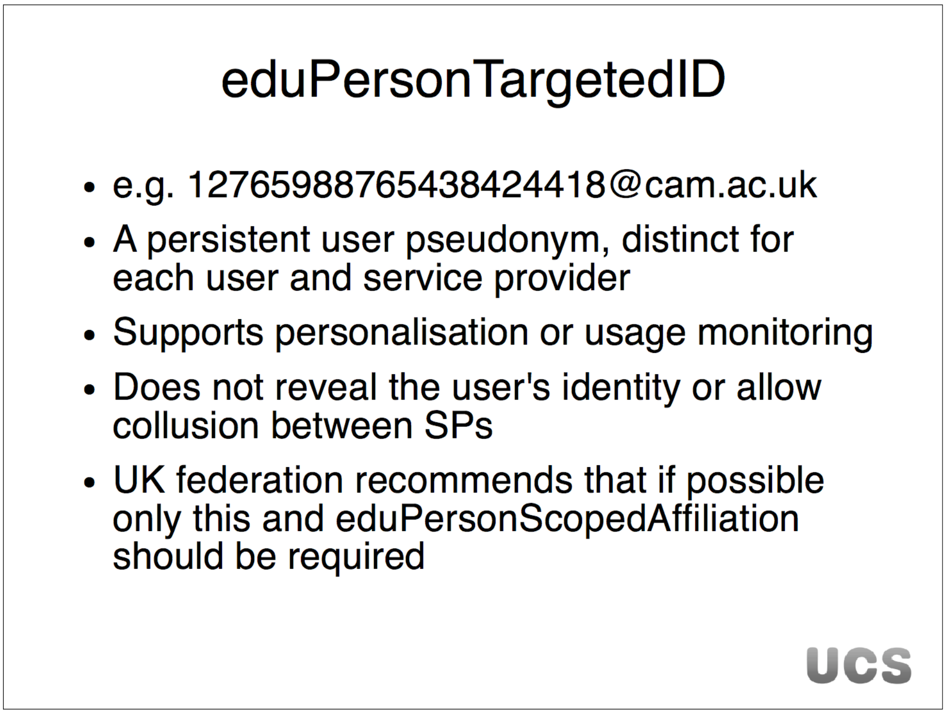

eduPersonTargetedID

- e.g. 12765988765438424418@cam.ac.uk
- A persistent user pseudonym, distinct for each user and service provider
- Supports personalisation or usage monitoring
- Does not reveal the user's identity or allow collusion between SPs
- UK federation recommends that if possible only this and eduPersonScopedAffiliation should be required

The eduPersonTargetedID (ePTID) attribute provides a persistent user pseudonym, which is distinct for each user and service provider. As such it can be used to support functions such as personalisation or usage profiling in a way that does not reveal the user's identity or allow collusion between SPs.

The Internet2 IdP software can generate ePTID on the fly by hashing the identity of the user and the SP with a secret. Alternatively, and perhaps better, ePTIDs can be generated when first required and then stored keyed to the user and SP identity. This latter approach allows ePTIDs to be changed if their anonymity is compromised, or to be maintained across a change to the representation of user or SP identity.

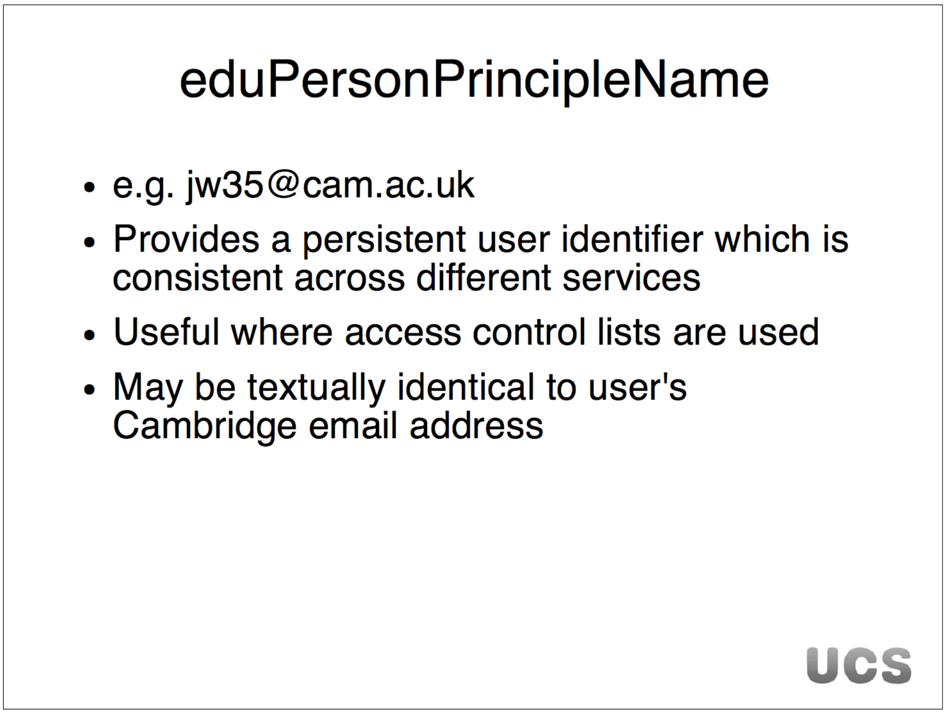

eduPersonPrincipleName

- e.g. jw35@cam.ac.uk
- Provides a persistent user identifier which is consistent across different services
- Useful where access control lists are used
- May be textually identical to user's Cambridge email address

This attribute is used where a persistent, unique user identifier is required that is explicitly consistent across different services. It often corresponds to the user's local single sign-on (SSO) name. It is especially useful where the identity has to be communicated out of band, perhaps by phone or email, or where access control lists are used. To support these uses there is an argument for constructing eduPersonPrincipleName (ePPN) from values with which users are already familiar.

UK Federation guidelines lead to a ePPN for someone at Cambridge having the form fjc55@cam.ac.uk. While textually identical to an @cam mail address, it is important to understand that this is neither an email address nor a Kerberos identity. This is likely to lead to user confusion.

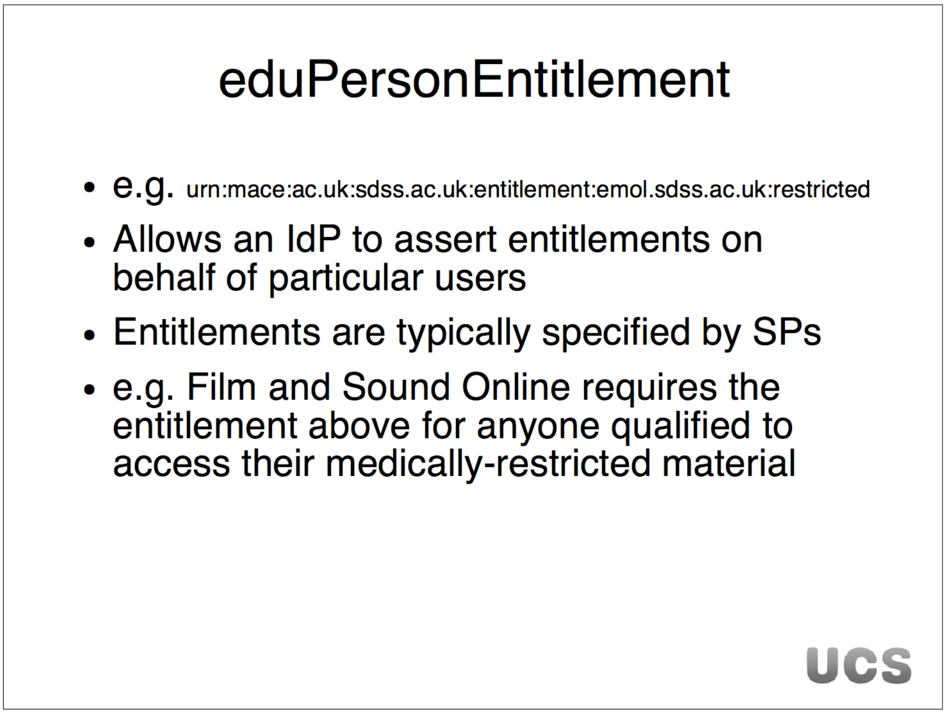

eduPersonEntitlement

- e.g. urn:mace:ac.uk:sdss.ac.uk:entitlement:emol.sdss.ac.uk:restricted
- Allows an IdP to assert entitlements on behalf of particular users
- Entitlements are typically specified by SPs
- e.g. Film and Sound Online requires the entitlement above for anyone qualified to access their medically-restricted material

The eduPersonEntitlement (ePE) attribute enables an organisation to assert that a user satisfies an additional set of specific conditions that apply for access to a particular resource. A user may possess different values of the eduPersonEntitlement attribute relevant to different resources.

This provides an escape mechanism allowing an IdP to assert one or more entitlements, typically specified by an SP, on behalf of particular IdP users. An example use case would be asserting that a particular user is entitled to access a particular resource under the terms of the relevant licence.

Values for ePE have the form of Uniform Resource Identifiers (URIs), most frequently using the 'http' or 'urn' schemes. In the case of a value using the 'http' scheme, the UK Federation recommends but does not require that the value resolve to a document giving the definition of the value.

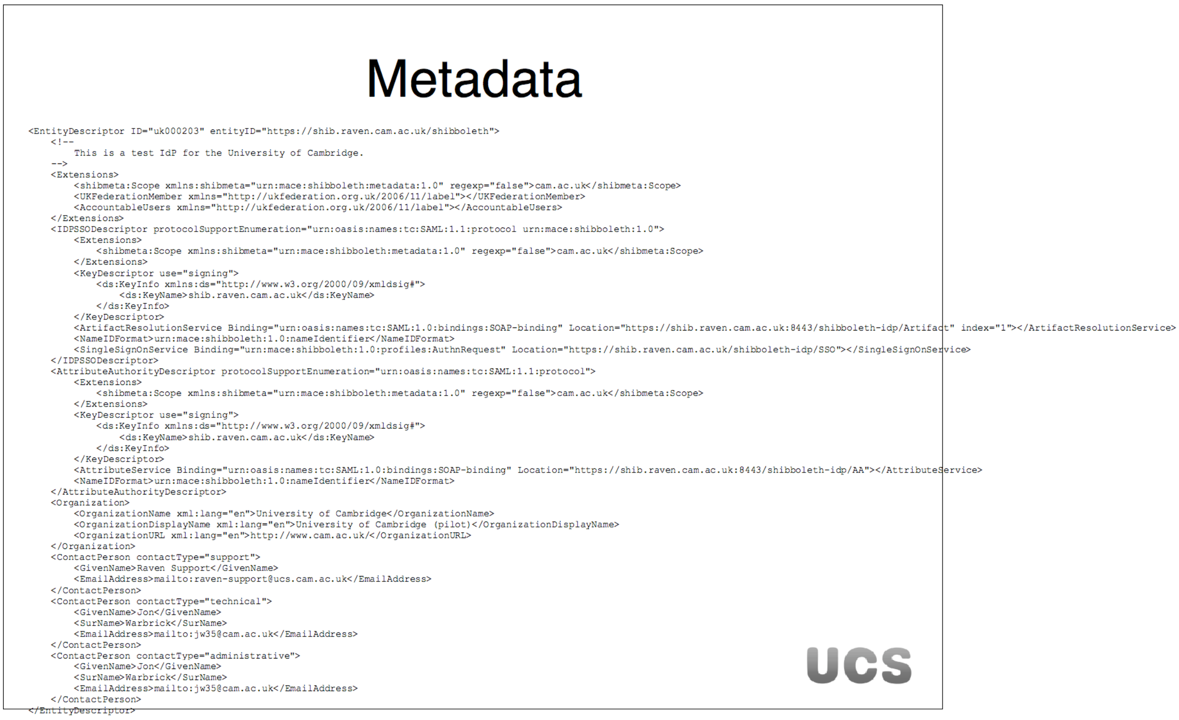

Metadata

```
<EntityDescriptor ID="uk000203" entityID="https://shib.raven.cam.ac.uk/shibboleth">
    <!--
        This is a test IdP for the University of Cambridge.
    -->
    <Extensions>
        <shibmeta:Scope xmlns:shibmeta="urn:mace:shibboleth:metadata:1.0" regexp="false">cam.ac.uk</shibmeta:Scope>
        <UKFederationMember xmlns="http://ukfederation.org.uk/2006/11/label"></UKFederationMember>
        <AccountableUsers xmlns="http://ukfederation.org.uk/2006/11/label"></AccountableUsers>
    </Extensions>
    <IDPSSODescriptor protocolSupportEnumeration="urn:oasis:names:tc:SAML:1.1:protocol urn:mace:shibboleth:1.0">
        <Extensions>
            <shibmeta:Scope xmlns:shibmeta="urn:mace:shibboleth:metadata:1.0" regexp="false">cam.ac.uk</shibmeta:Scope>
        </Extensions>
        <KeyDescriptor use="signing">
            <ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
                <ds:KeyName>shib.raven.cam.ac.uk</ds:KeyName>
            </ds:KeyInfo>
        </KeyDescriptor>
        <ArtifactResolutionService Binding="urn:oasis:names:tc:SAML:1.0:bindings:SOAP-binding" Location="https://shib.raven.cam.ac.uk:8443/shibboleth-idp/Artifact" index="1"></ArtifactResolutionService>
        <NameIDFormat>urn:mace:shibboleth:1.0:nameIdentifier</NameIDFormat>
        <SingleSignOnService Binding="urn:mace:shibboleth:1.0:profiles:AuthnRequest" Location="https://shib.raven.cam.ac.uk/shibboleth-idp/SSO"></SingleSignOnService>
    </IDPSSODescriptor>
    <AttributeAuthorityDescriptor protocolSupportEnumeration="urn:oasis:names:tc:SAML:1.1:protocol">
        <Extensions>
            <shibmeta:Scope xmlns:shibmeta="urn:mace:shibboleth:metadata:1.0" regexp="false">cam.ac.uk</shibmeta:Scope>
        </Extensions>
        <KeyDescriptor use="signing">
            <ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
                <ds:KeyName>shib.raven.cam.ac.uk</ds:KeyName>
            </ds:KeyInfo>
        </KeyDescriptor>
        <AttributeService Binding="urn:oasis:names:tc:SAML:1.0:bindings:SOAP-binding" Location="https://shib.raven.cam.ac.uk:8443/shibboleth-idp/AA"></AttributeService>
        <NameIDFormat>urn:mace:shibboleth:1.0:nameIdentifier</NameIDFormat>
    </AttributeAuthorityDescriptor>
    <Organization>
        <OrganizationName xml:lang="en">University of Cambridge</OrganizationName>
        <OrganizationDisplayName xml:lang="en">University of Cambridge (pilot)</OrganizationDisplayName>
        <OrganizationURL xml:lang="en">http://www.cam.ac.uk/</OrganizationURL>
    </Organization>
    <ContactPerson contactType="support">
        <GivenName>Raven Support</GivenName>
        <EmailAddress>mailto:raven-support@ucs.cam.ac.uk</EmailAddress>
    </ContactPerson>
    <ContactPerson contactType="technical">
        <GivenName>Jon</GivenName>
        <SurName>Warbrick</SurName>
        <EmailAddress>mailto:jw35@cam.ac.uk</EmailAddress>
    </ContactPerson>
    <ContactPerson contactType="administrative">
        <GivenName>Jon</GivenName>
        <SurName>Warbrick</SurName>
        <EmailAddress>mailto:jw35@cam.ac.uk</EmailAddress>
    </ContactPerson>
</EntityDescriptor>
```

UCS

Another function of a federation is to distribute *metadata* about all the SPs and IdPs operated by their members. This lets members identify other members, defines how their various services can be found, etc.

This metadata is distributed in XML, and signed by the federation operator to prevent tampering. This slide contains (part of) the metadata about the University's IdP.

Why should we bother?

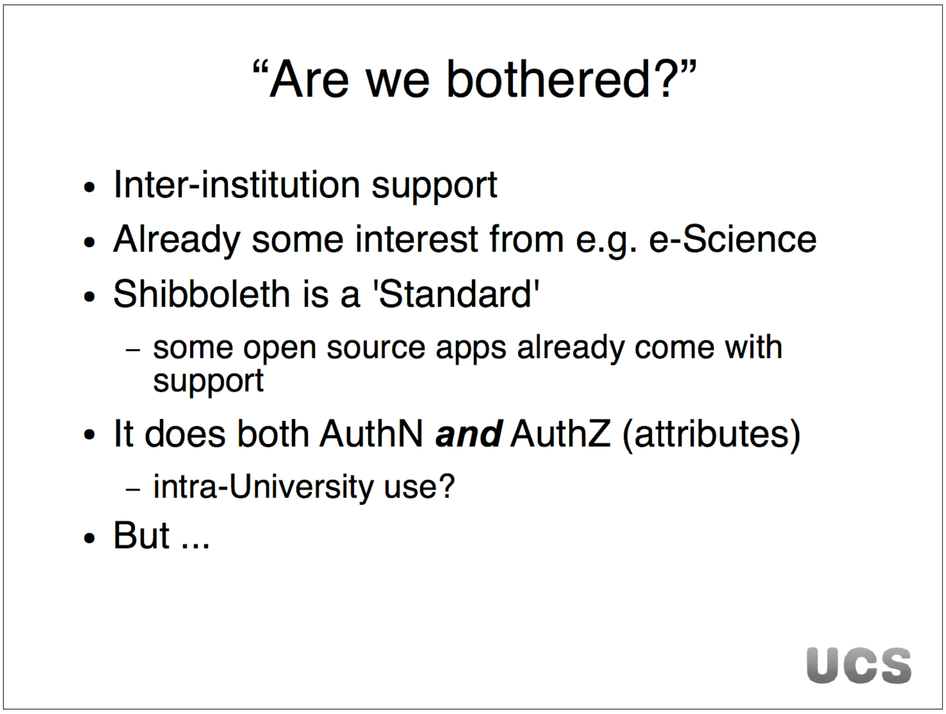

“Are we bothered?”

- Inter-institution support
- Already some interest from e.g. e-Science
- Shibboleth is a 'Standard'
  - some open source apps already come with support
- It does both AuthN ***and*** AuthZ (attributes)
  - intra-University use?
- But ...

UCS

The ability to identify people from places other than Cambridge is a much-requested feature for Raven which Shibboleth supplies 'out of the box'. There is already some interest in Shibboleth from areas such as the e-Science community who already have to manage users from multiple institutions.

The fact that Shibboleth is a standard makes it likely that third-party software will support it directly – current versions of Media-wiki, Moodle, dSpace, MyProxy, and uPortal all claim some level of Shibboleth support.

Delivering attribute information along with authentication is another feature commonly requested of Raven that Shibboleth addresses, so it's likely that Shibboleth will eventually be used to support intra-University applications.

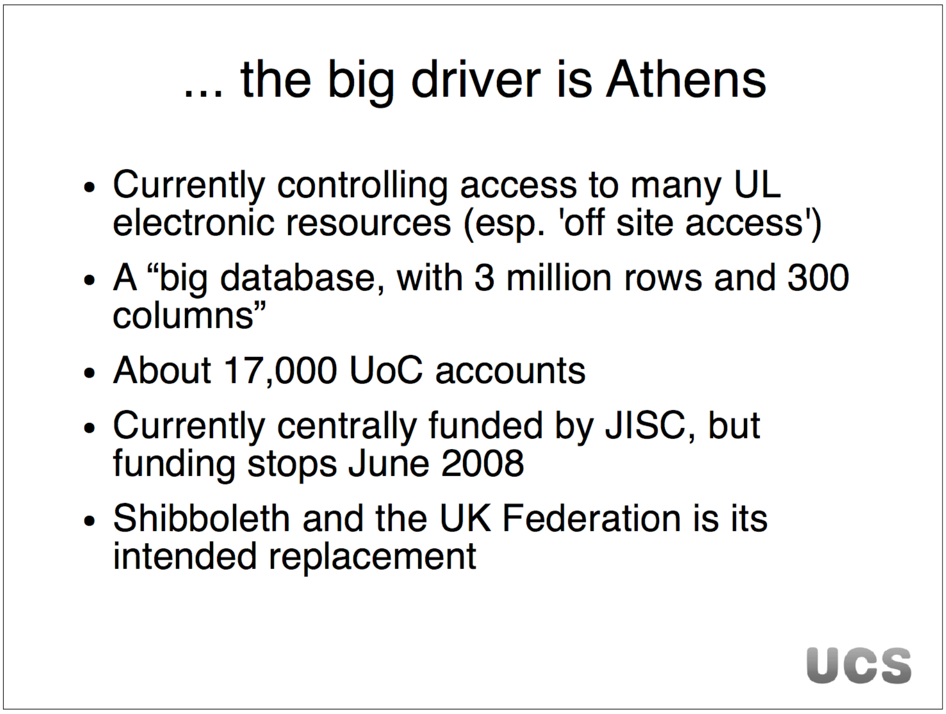

... the big driver is Athens

- Currently controlling access to many UL electronic resources (esp. 'off site access')
- A “big database, with 3 million rows and 300 columns”
- About 17,000 UoC accounts
- Currently centrally funded by JISC, but funding stops June 2008
- Shibboleth and the UK Federation is its intended replacement

Athens is a national service run by EduServ under contract for JISC which is used to authenticate and authorise access to many library electronic resources. Athens is in effect a national userid and password storage and verification system. In the University, Athens authentication is currently managed by the University Library. We have approximately 17,000 Athens users.

The University's use of Athens has always been centrally funded by JISC, but this funding ceases in June 2008. EduServ plan to continue to operate the service on a commercial basis for the foreseeable future. JISCs official strategy is to replace Athens by Shibboleth.

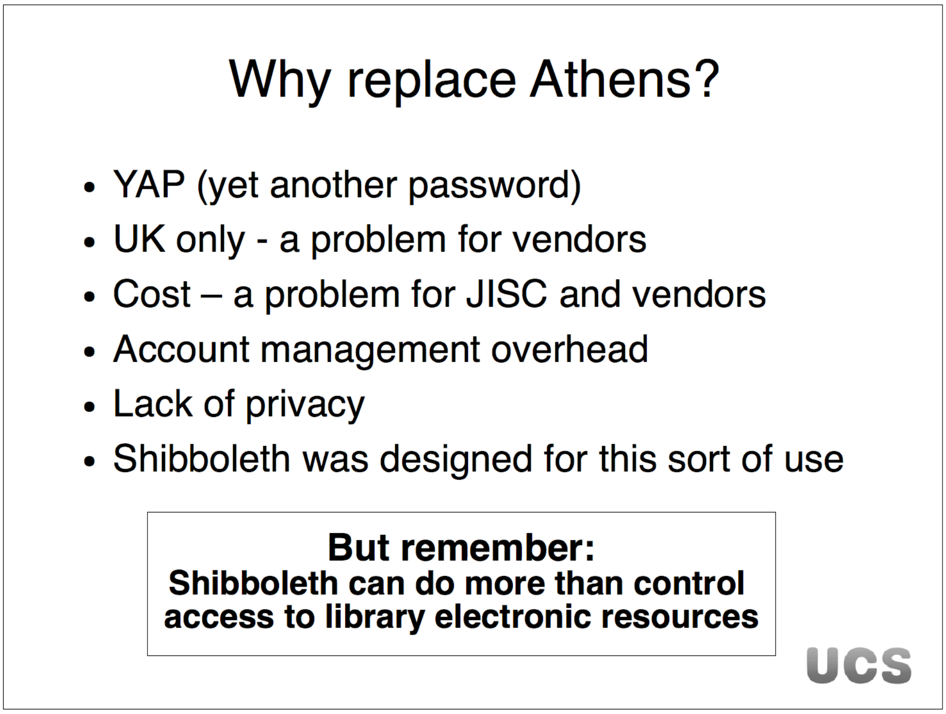

Why replace Athens?

- YAP (yet another password)
- UK only - a problem for vendors
- Cost – a problem for JISC and vendors
- Account management overhead
- Lack of privacy
- Shibboleth was designed for this sort of use

But remember:
Shibboleth can do more than control access to library electronic resources

UCS

There are a number of reasons why replacing Athens by something like Shibboleth is a good idea. Perhaps the most important is that it reduces by one the number of distinct userids and passwords our users have to remember. The fact that we will no longer have to manage the Athens id/password set is also worthwhile.

Athens is a problem for SPs too – it is UK-specific and so requires implementation work on their side. They, and JISC, also have to pay a significant sum to EduServ for the Athens service.

Athens lacks the privacy-preserving options of Shibboleth.

Shibboleth was originally designed with access control for electronic resources in mind, but it is important to remember that this is far from the only thing it can do.

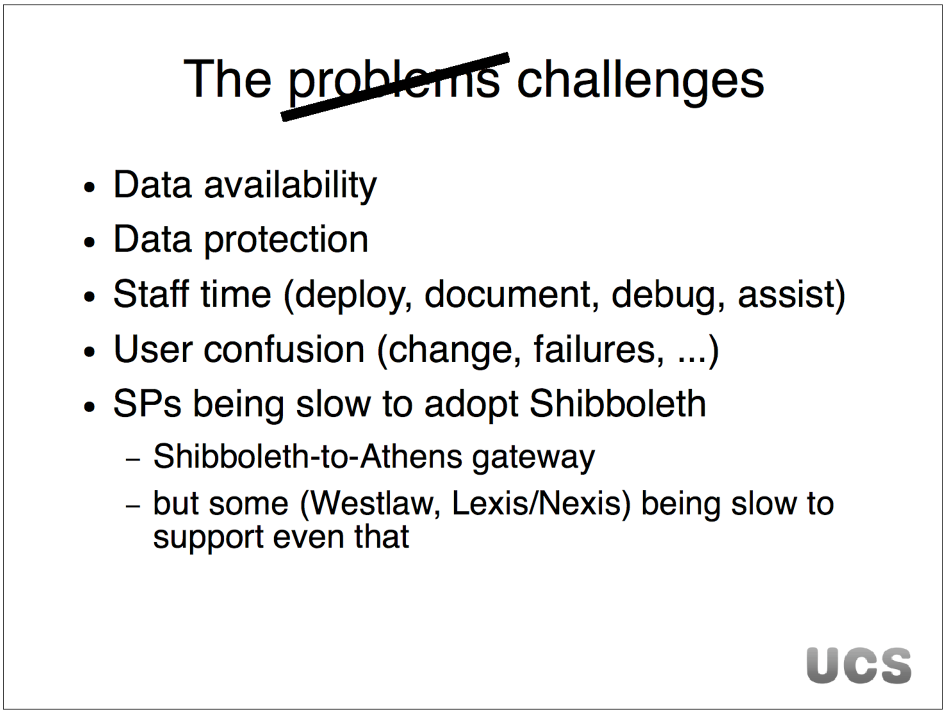

The ~~problems~~ challenges

- Data availability
- Data protection
- Staff time (deploy, document, debug, assist)
- User confusion (change, failures, ...)
- SPs being slow to adopt Shibboleth
  - Shibboleth-to-Athens gateway
  - but some (Westlaw, Lexis/Nexis) being slow to support even that

Implementing something as complex as Shibboleth is never going to be entirely simple. While we are better off than many institutions (because in particular we have Raven and lookup) there are still gaps in the data we need - e.g. reliable lists of College staff. In the UK we also have the requirements of the Data Protection Act to comply with, something which doesn't directly affect the Internet2 developers in the US.

Deploying Shibboleth is taking a lot of staff time in the CS, the UL and elsewhere. It's a new technology and it is bound to go wrong in new and unexpected ways. Users who have been happily using Athens for many years will not be pleased by the change, especially since it may result in their existing customisations of many services getting lost.

Existing Athens suppliers are being slow to adopt Shibboleth. Few suppliers currently support it directly, though many more can be accessed via a Shibboleth-to-Athens gateway provided by EduServ. However some don't even work through the gateway, making it problematic to drop Athens complelty.

When is it all going to happen?

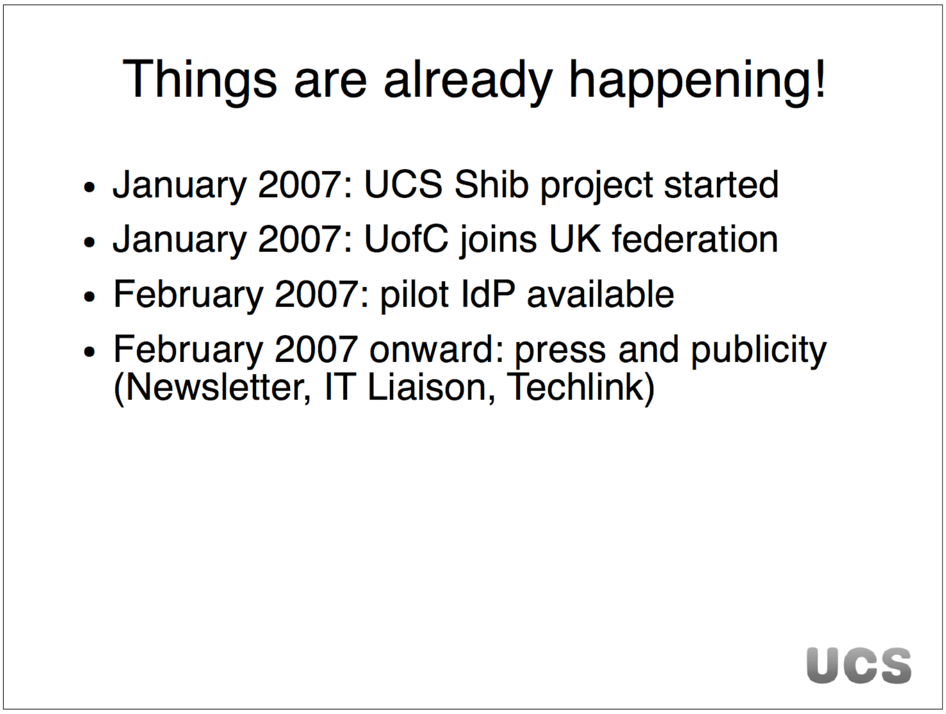

Things are already happening!

- January 2007: UCS Shib project started
- January 2007: UofC joins UK federation
- February 2007: pilot IdP available
- February 2007 onward: press and publicity (Newsletter, IT Liaison, Techlink)

UCS

Work by the CS and UL on deploying Shibboleth started in January 2007. We have a demonstration-quality IdP in place and have successfully joined the UK Federation. Various publicity events have been arranged, including this talk. The work undertaken so far has put us in a good position to move forward to provide a production service.

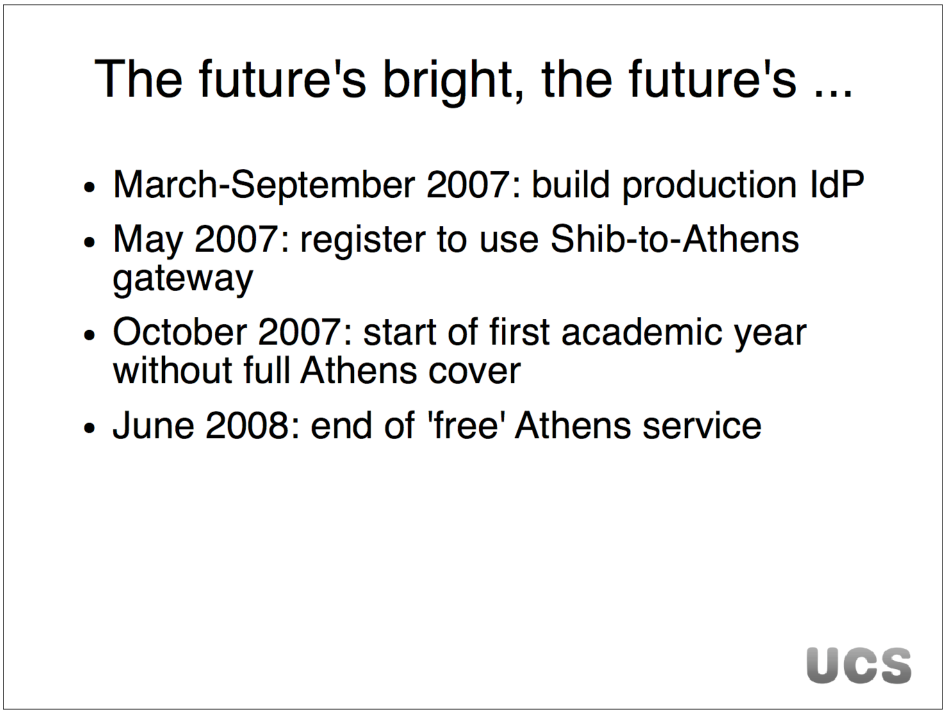

The future's bright, the future's ...

- March-September 2007: build production IdP
- May 2007: register to use Shib-to-Athens gateway
- October 2007: start of first academic year without full Athens cover
- June 2008: end of 'free' Athens service

The next job is to transition the current demonstration-quality IdP into something that can support production use. This work is under way and is expected to be complete in time for the start of the 2007/08 academic year.

In parallel with this, the UL is registering to use the Shibboleth-to-Athens gateway. Once this is available, they will be able to evaluate using pure-Shibboleth to control access to the electronic resources that they manage, and this will inform their future plans. One option would be to use pure Shibboleth for new arrivals from October 2007 and to transition existing Athens users to Shibboleth before the June 2008 cut-off date. Alternatively if the technology proves insufficiently mature there remains the option of paying to use Athens for one or more years until the situation improves.

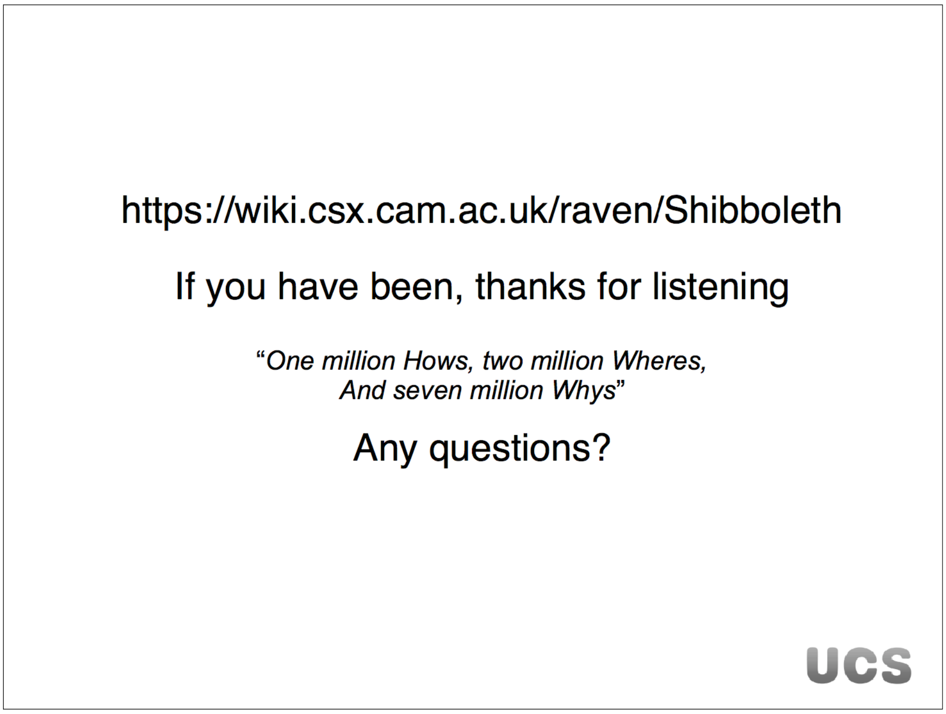

https://wiki.csx.cam.ac.uk/raven/Shibboleth

If you have been, thanks for listening

"One million Hows, two million Wheres,
And seven million Whys"

Any questions?

Further information about the Shibboleth project, including links to related material and a copy of this presentation, is available at

https://wiki.csx.cam.ac.uk/raven/Shibboleth

Questions and comments can be addressed to

raven-support@ucs.cam.ac.uk