

Internet2 Shibboleth Project ... and the UW

RL "Bob" Morgan, C&C
UW Computing Quarterly Support Meeting
March 17, 2004

Topics

Internet2 Middleware Initiative

Shibboleth basics

Current status and next steps

Federations and InCommon

Shibboleth, UW, pubcookie, and all that

Internet2 Middleware Initiative

- develop best practice, recipes, experiments, new tech
 - started 1999
- focus on security/directory infrastructure
 - i.e., finding stuff, and hiding stuff
- Working Groups on many topics
 - Directory: eduPerson, Groups
 - Security: PKI, Webiso (aka web SSO), Shibboleth
 - App integration: Instant Messaging, Video conferencing, P2P, medical, network management, ...
 - Naming/identifiers: URNs, OIDs
- <http://middleware.internet2.edu/>

MACE (Middleware Architecture Committee for Education)

Purpose - steering committee for I2MI: to provide advice, foster standards, etc. on key technical issues for core middleware within higher education

Membership - Bob Morgan (UW) Chair, Tom Barton (Chicago), Scott Cantor (Ohio State), Steven Carmody (Brown), Michael Gettes (Duke), Keith Hazelton (Wisconsin), Paul Hill (MIT), Jim Jokl (Virginia), Mark Poepping (CMU), Bruce Vincent (Stanford), David Wasley (California), Von Welch (Grid)

European members - Brian Gilmore (Edinburgh), Ton Verschuren (Netherlands), Diego Lopez (Spain)

Works via conference calls, emails, occasional serendipitous in-person meetings...



The National Science Foundation Middleware Initiative (NMI)

NSF Program to develop and deploy common
middleware infrastructures

Two major themes

- Scientific computing and data environments (ala Grids)

- Common campus and inter-institutional middleware
infrastructure (ala Internet2/EDUCAUSE/SURA work)


“Random acts of middleware” projects to
complement major systems integration work

Issues periodic NMI releases of software, services,
architectures, objectclasses and best practices – R4
due out the end of the year

A Map of Middleware Land

NMI-

Enabling Integral Science,
Engineering, Research,
and Education.



Internet2 Middleware Principles

“Federated” model: independent entities provide core services, central services facilitate discovery, trust

Consistent directory infrastructure within R&E

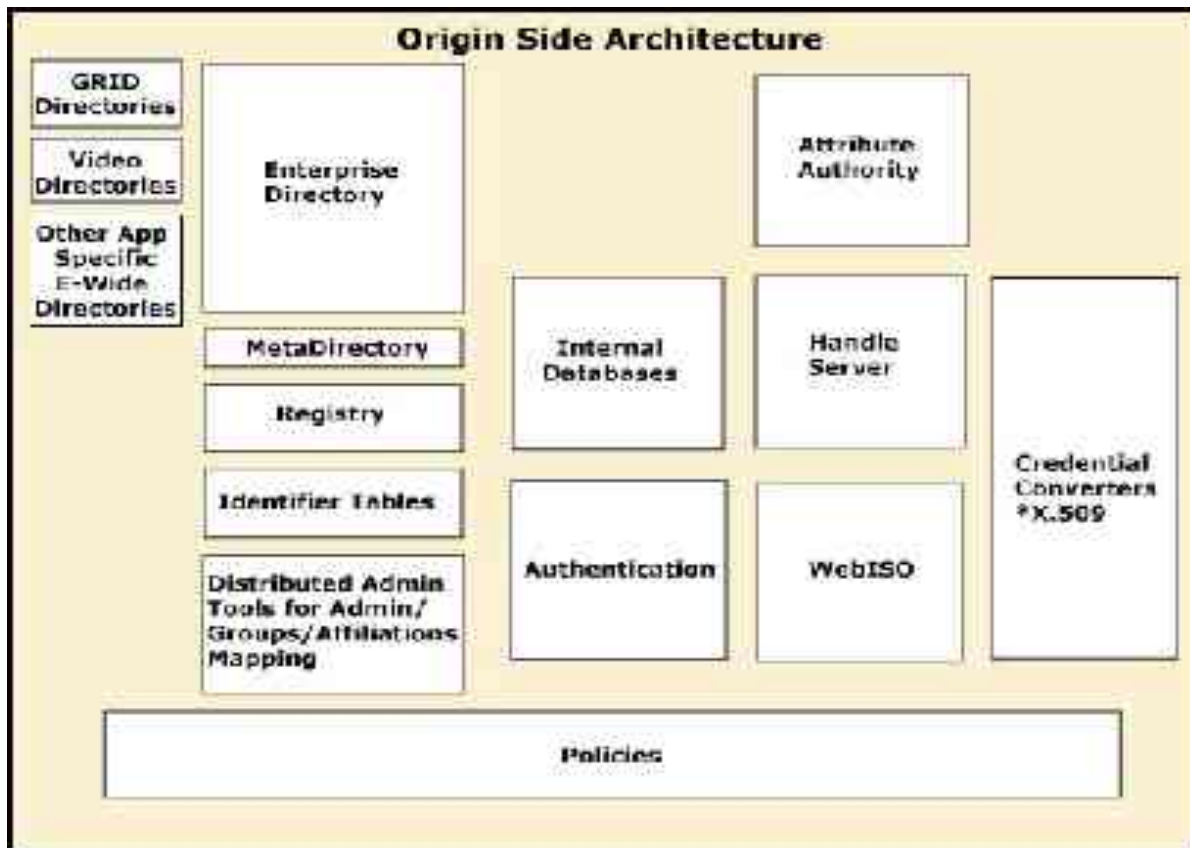
Good security without degrading privacy

Leverage campus expertise and build rough consensus

Influence the marketplace; develop where necessary

Support for heterogeneity, open standards, open source

Campus Core Middleware



Interrealm and intrarealm

“Intrarealm”: describes the services within an enterprise, such as a university. Middleware services assume commonalities and trust.

“Interrealm”: describes relationships between autonomous systems or enterprises. Service agreements based on contracts.

Universities have many pockets of semi-autonomy (colleges, medical schools and hospitals, athletic departments) may best be viewed as interrealm ...

•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 a continuum from interrealm groups (no real resource management, few defined roles) to real organizations (primary identity/authentication providers)

A required cross-stitch on the enterprise trust fabric

Stay tuned...

Shibboleth: starting points

- Web-centric
 - most new info resources provided via web, especially inter-realm
 - campuses often deploying web-based central sign-on
- Focus on key inter-realm deployment scenarios ...
- Authorization-centric
 - apps want “sign-on” to do access control / authorization
 - often replacing non-signon-based authorization
 - focus on attribute-based authz for privacy support
- Industry standards work on web signon
 - i.e., OASIS Security Assertion Markup Language (SAML)
 - benefit from clue, vendor support, standards process

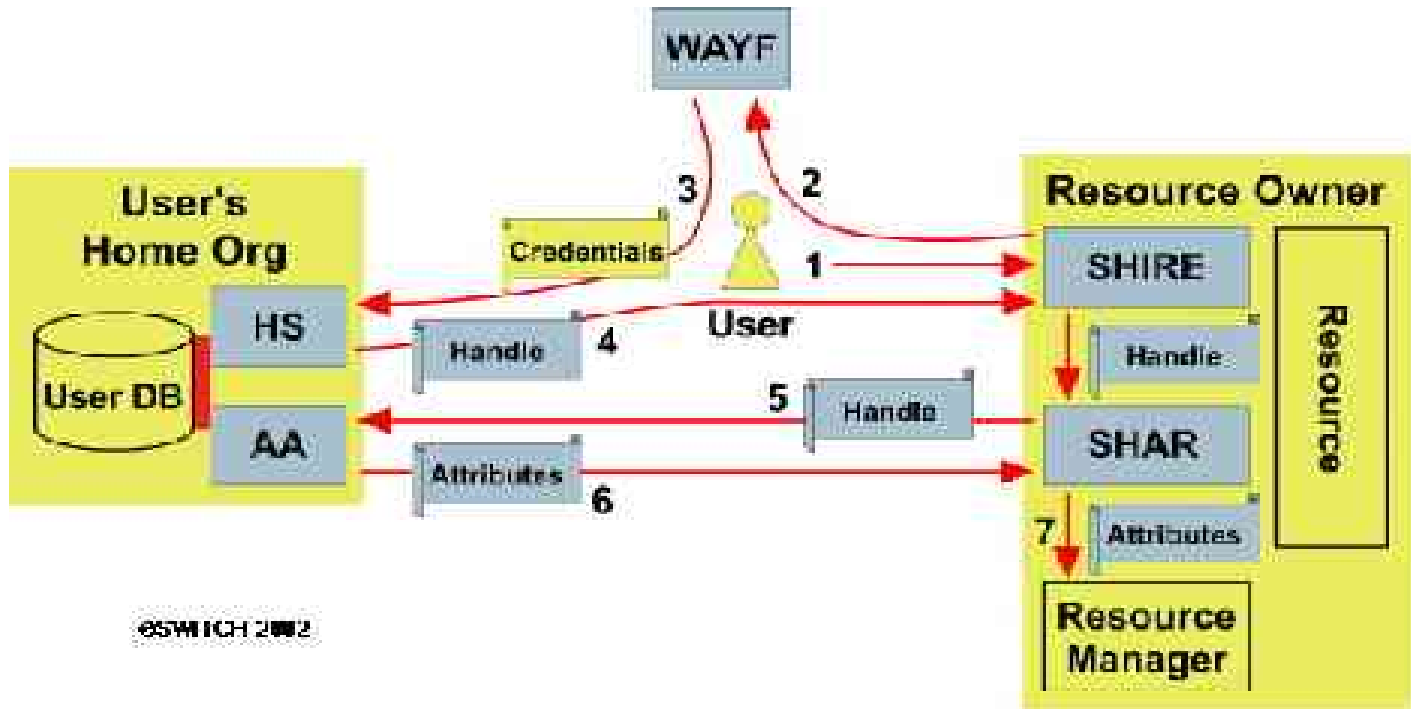
Three Shibboleth Scenarios

- (1) Member of campus community accessing licensed resource
Privacy-protection (aka anonymity) important
Integration with contract/licensing process
- (2) Member of group (e.g., course) accessing remote resource
Privacy-protection important in many cases
Integration with group processes (e.g. course mgmt)
- (3) Member of a workgroup accessing controlled resources
Based on unique identifiers (i.e., userids)

Taken individually, each may be solved somewhat easily.

Together, challenges of authentication, authorization, privacy protection, common services.

• Shibboleth Architecture



Shibboleth components

- Protocol: SAML web browser profile, attribute query/resp
 - XML-based formats for “security assertions” and protocol
 - http(s) POST to submit SAML authentication assertion to target
 - SAML/SOAP/https to get attributes from origin Attribute Authority
 - profiled for Shib: naming, keys, attributes, ...
- Target
 - in C/C++ for Apache/UNIX, IIS
 - Java target under development
 - site configuration data describes origins, names, keys
 - “attribute acceptance” says which origins can say what
 - standard .htaccess / env-vars for per-app configuration

Shibboleth components: origin

- Shibboleth Handle Service
 - aka “authentication authority”
 - translates site-local signon into SAML authn assertion
 - generates short-lived identifiers for use with Attr Authority
 - interacts with browser to provide POSTable assertion
- Attribute Authority
 - responds to SAML queries from targets
 - retrieves attributes from local store (eg LDAP directory) or generate attributes based on rules
 - filters based on “attribute release policies” per user/group
 - management tools: ARPs, other

more Shibboleth components

- WAYF (
 - aka “origin discovery”
 - how do users get from target back to origin to do authn
 - can be within target service, or provided by third party
- policy agreements
 - how does target know about origin security methods?
 - e.g., are all users equally well-authenticated?
 - privacy policies imposed on targets
- attribute definitions
- underlying security arrangements
 - rely on PKI, but many flavors to choose from ...

Shibboleth Milestones

Project formation - Feb 2000

Process - began late summer 2000 with bi-weekly calls to develop scenario, requirements and architecture.

Linkages to OASIS / SAML established Dec 2000

Architecture and protocol completion - Aug 2001

Coding start - Nov 2001

Alpha-1 release – April 24, 2002

OpenSAML release – July 15, 2002

v1.0 April 2003

v1.1 July 2003

v1.2 April 2004

• Shibboleth Status

40-50 “origin” campuses in trial group

6 or so major resource providers, other products as targets

Relatively straightforward to install, provided there is good web services understanding and middleware infrastructure (authentication, directories, webISO, etc.).

Work underway on some of the essential management tools such as attribute release managers, target resource management, etc.

Likely to coexist well with Liberty Alliance and may work within the Web Services framework from Microsoft/IBM.

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.)

•Federations

Associations of enterprises that come together to exchange information about their users and resources in order to enable collaborations and transactions

Built on the premise of

Initially “Authenticate locally, act globally”

Now, “Enroll and authenticate and attribute locally, act federally.”

Federation provides only modest operational support and consistency in how members communicate with each other

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.

Over time, this will all change...

•Requirements for federations

Federation administration and operations

Federating software

- Exchange assertions

- Link and unlink identities

Federation data schema

Federation privacy and security requirements

• Shibboleth-based federations

Today (or very soon)

InQueue

InCommon

SWITCH

NSDL

Future

State networks

Medical networks

Financial aid networks

Life-long learning communities

InQueue

The “holding pond”

A persistent federation with “passing-through” membership of sites coming up to speed

Operational today. Can apply for membership via <http://shibboleth.internet2.edu/> InQueue Federation guidelines

Requires eduPerson attributes

Operated by Internet2; open to organizations using Shibboleth in an Research&Ed setting, or not...

Fees and service profile to be established shortly: cost-recovery basis

Major targets

Campuses that are also origins, wanting to share campus-based content

Content providers – EBSCO, OCLC, JSTOR, Elsevier, Napster (PSU), etc

Next round prospectives: Ovid, CSA, Gale Group, ISI, Proquest, Lexis-Nexis

Learning Management Systems – WebCT, Blackboard, OKI, etc

Outsourced Service Providers – purchasing systems, dormitory management companies, course management services, etc.

InCommon federation

Permanent managed federation for the R&E US sector
oversight/management by campus CIOs

Operated by Internet2

open to .edu-qualified sites and business partners
sold as “brand” users will see at participating targets

Security requirements:

Initially, enterprises post local I/A and basic business rules for
assignment of eduPersonAffiliation values

Likely to progress towards standardized levels of authn

may expand to federate other technologies

IM, video, spam(?), PKI, etc.

•Shib academic SIG

Lots of interesting design issues for use of Shib, e.g

- Passing attributes during deep-linked text

- Handling meta-search engines

- Managing persistent identifiers where needed

- Dealing with proxies in a semi-Shibbed world

The issues so far have all been solvable; the challenge is in picking the right solution.

Subscribe and participate via the I2 listserv at <http://www.internet2.edu/about/lists.html> (sigh, soon to be Shibboleth-enabled...)

Computing industry activities

- much buzz about “federated identity”
- many products support SAML
- Liberty Alliance
 - profile of SAML for big-provider-centric federation and services
 - Liberty work being folded into SAML 2.0
- Microsoft
 - vast amount of work on “Web Services framework”
 - highly modularized set of services ...
 - WS-Federation is competitor to SAML browser standard
 - convergence possible ...
- real federation, and federations, happening in corp space

Shib and UW

- operating origin in “pilot” mode
 - production planning starting, in context of pubcookie project
 - operation of attribute authority raises interesting questions
- so, what about pubcookie?
 - Shib is complementary, now and for quite a while
 - move to standard SAML formats possible
 - Shib can't do everything pubcookie can, yet ...
 - weblgin service, securid mode, logout, forced-authn, performance?
- pubcookie vs Shib target, which to use ...
 - Shib may be better in many vendor scenarios
 - e.g., higher-ed application service providers
 - many now implementing multiple web signon schemes ...

Attributes

- pubcookie just supplies userid (aka UW NetID)
- Shib supplies attributes as part of normal ops
- many sites interested in user attributes
 - eg “faculty”, “course xyz”
- within UW, probably better provided by LDAP lookup
 - working on “Enterprise Directory” to provide this, real soon now

Shib opportunities

- UW resources provided to other sites ...
 - Research Channel, Digital Well
 - Catalyst
 - working on Shib target support to work with other colleges/univs
 - kind of like Y2K: assumptions about “netids”
 - reduce likelihood of UWNNetIDs for non-UW organizations?
- UW users using other sites ...
 - e.g. WebAssign
 - other ASPs
- your department's collaborative project?
 - will the other side have Shib origin deployed?

Conclusion

- Shibboleth stable, in use
- Operational federations coming together
- Lots of additional layers to build
- Lots of integration left to do with
 - content providers
 - institutional user management systems
 - vertical target apps
- Deployment at UW starting to happen
 - Integration into our infrastructure will be ongoing topic ...