The Rise of Collaborative Tools: The Golden Age of Plywood

Ken Klingenstein
Project Director, Internet2 Middleware Initiative
Chief Technologist, University of Colorado at Boulder

Topics

- *Background and acknowledgments*
- *Generalisms and running water*
- *A few basic concepts*
- *The plumbing*
  - Directories
  - Shibboleth
  - Federations
- *The tools*
- *Issues and opportunities for OCLC*

MACE (Middleware Architecture Committee for Education)

- *Purpose - to provide advice, create experiments, foster standards, etc. on key technical issues for core middleware within higher education*
- *Membership - Bob Morgan (UW) Chair, Scott Cantor (Ohio State), Steven Carmody (Brown), Michael Gettes (Georgetown), 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)*
- *Creates working groups in major areas, including directories, interrealm access control, PKI, medical issues, etc.*
- *Works via conference calls, emails, occasional serendipitous in-person meetings...*

The National Science Foundation Middleware Initiative (NMI)

- *NSF award for integrators to*
  - Grids (NCSA, UCSD, University of Chicago, USC/ ISI, and University of Wisconsin)
  - Internet2, EDUCAUSE, and SURA
- *Build on the successes of the Globus project and the Internet2/MACE initiative*
- *Multi-Year Effort*
- *A practical (deployment) activity that necessitates some research and much development*
- *Separate awards to academic pure research “throw it long” components*
- *Issues periodic NMI releases of software, services, architectures, objectclasses and best practices – R4 due out the end of the year*

Making it happen

- *Much as at the network layer, plumb a ubiquitous common, persistent and robust core middleware infrastructure for the R&E community*
    - Foster effective and consistent campus implementations
    - Motivate institutional funding and deployment strategies
    - Solve the real world policy issues
    - Integrate key applications to leverage the infrastructure
    - Nurture open-source solutions
    - Address scaling issues for the user and enterprise
- *in support of inter-institutional and interrealm collaborations, provide tools and services (e.g. registries, bridge PKI components, root directories) as required*

The way we work

- *Members such as MIT, CMU, Wisconsin, Ohio State, Washington and many others provide their best staff as volunteers and part-time workers*
- *Corporations such as SUN and IBM contribute software, expertise, willingness to shape products, etc.*
- *Government agencies such as NSF, NIST and NIH contribute services, expertise, research*
- *Internet2 staff act as layclergy and flywheels.*

Internet2 Middleware:
Key Concepts

- *Use federated administration as the lever; have the enterprise broker most services (authentication, authorization, resource discovery, etc.) in inter-realm interactions*
- *Develop a consistent directory infrastructure within R&E*
- *Provide security while not degrading privacy.*
- *Foster interrealm trust fabrics*
- *Leverage campus expertise and build rough consensus*
- *Influence the marketplace; develop where necessary*

Interrealm and intrarealm

- *Interrealm describes the relationships between autonomous systems or enterprises.*
- *Intrarealm describes the services within an enterprise, such as a university or corporation. The services, such as authentication, authorization and directories, assume commonalities and trust.*
- *But of course, for most large universities, there are many pockets of semi-autonomy (colleges, medical schools and hospitals, athletic departments) and it may best be viewed as interrealm*
- *And, of course, in large companies with many wholly-owned but acquired subsidiaries, the lack of a common infrastructure makes their architectures interrealm.*

Upper and Core Middleware Land

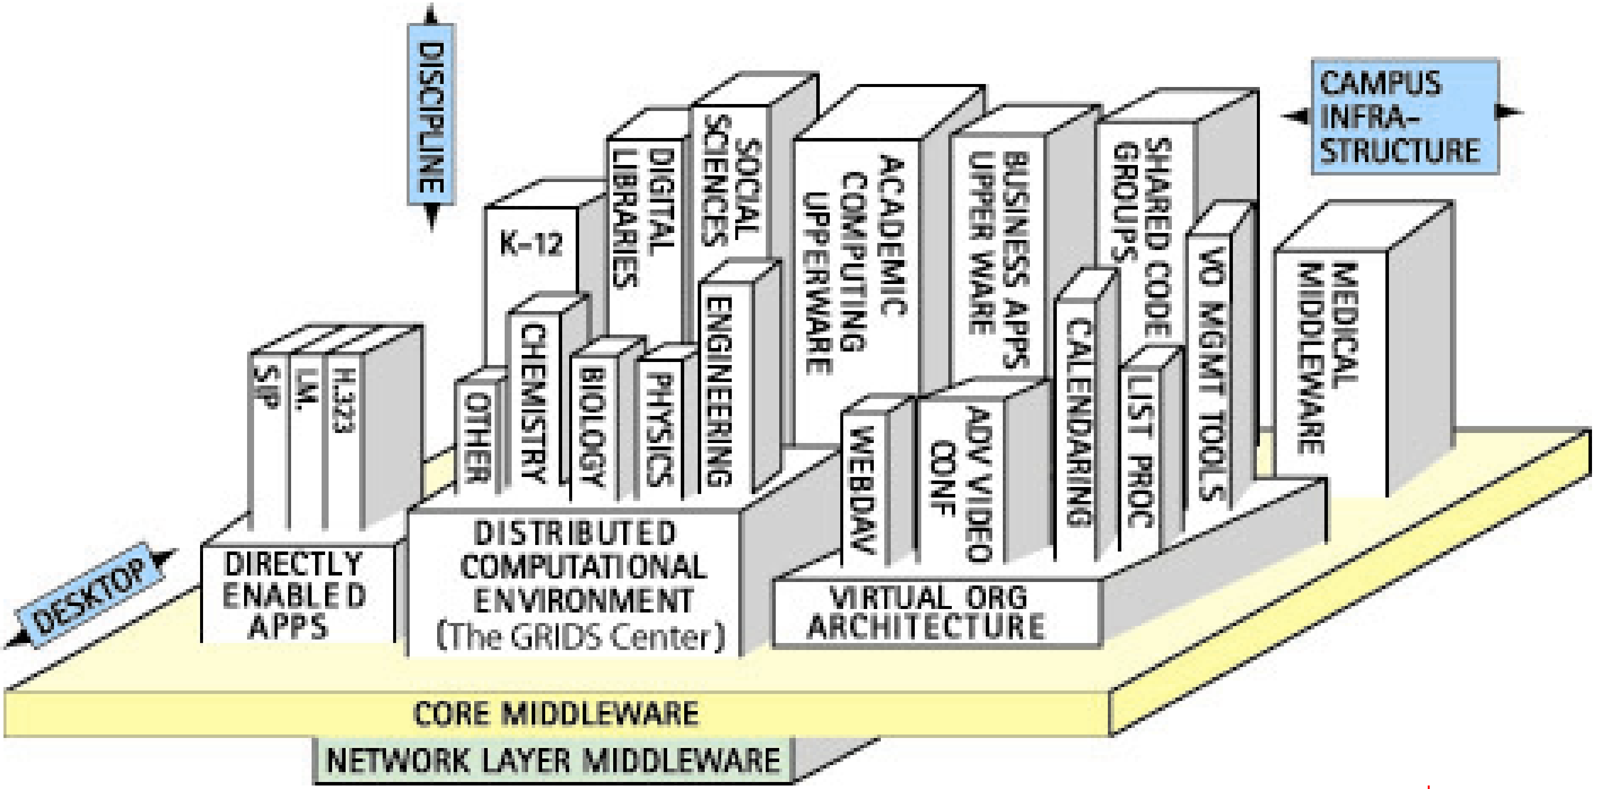

Core Middleware Scope

- ***Identity and Identifiers*** *– namespaces, identifier crosswalks, real world levels of assurance, etc.*
- ***Authentication*** *– campus technologies and policies, interrealm interoperability via PKI, Kerberos, etc.*
- ***Directories*** *– enterprise directory services architectures and tools, standard objectclasses, interrealm and registry services*
- ***Authorization*** *– permissions and access controls, delegation, privacy management, etc.*
- ***Integration Activities –*** *open management tools, use of virtual, federated and hierarchical organizations, enabling common applications with core middleware*

Campus Core Middleware Architecture: (Origin perspective)

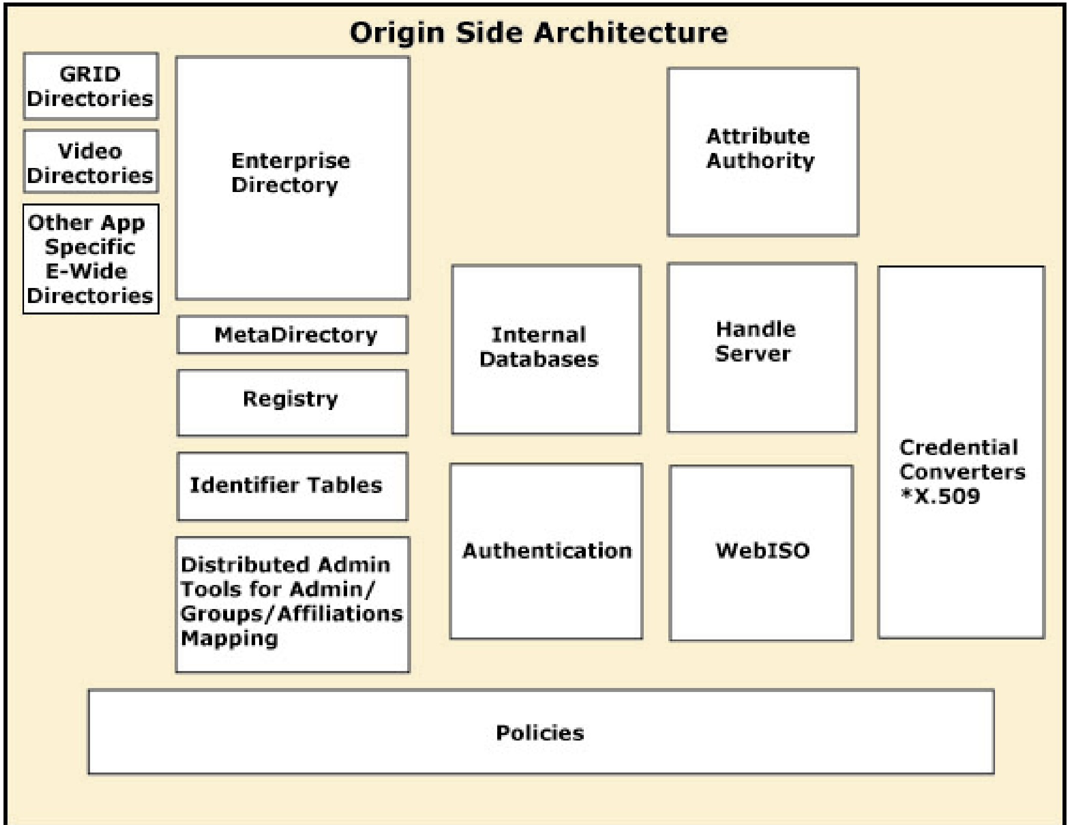

Origin Side Architecture

GRID Directories

Video Directories

Other App Specific E-Wide Directories

Enterprise Directory

MetaDirectory

Registry

Identifier Tables

Distributed Admin Tools for Admin/ Groups/Affiliations Mapping

Attribute Authority

Internal Databases

Handle Server

Credential Converters *X.509

Authentication

WebISO

Policies

Related upper middleware in Higher Ed

- *OKI*
- *IMS*
- *Grids*
- *Chandler*
- *Lionshare*

Generalisms and Running Water

- *OKI-IMS-Shibboleth-InCommon-Grids-andeverythingelse*
- *There is a game to be more general than another, but at some point the generalizations are too broad to be useful*
- *Users and deployers don’t care about generalizations or abstractions; they want running water and they want ¾ inch pipe to connect to ¾ inch valves*
- *Some things are abstractions (OKI, Grids, etc) and some are plumbing (Shib, Globus, etc.)*
- *We need to show how the plumbing is instances of the generalizations, and warn that other instances of the same generalizations may not interoperate, but should…*

Landmark Work

- *Convincing ourselves that we could do this and that it would make a difference…*
- *Consensus standards – eduPerson, eduOrg, commObject*
- *Best Practices and Deployment Strategies – LDAP Recipe, Group Management, Metadirectories*
- *Tools – KX.509, LDAP Analyzer, LOOK*
- *Software systems – OpenSAML, Shibboleth*

A method to the madness

- *First develop campus directories – the LDAP recipe*
- *Then develop an external-facing collaborative-oriented set of attributes to put in those directories – eduPerson, eduOrg*
- *Then develop an privacy-preserving, secure inter-realm transport for those attributes – Shibboleth*
- *Now develop methods to trust the attributes that are transported – federations (InCommon, etc)*
- *Soon apply to services beyond the web – videoconferencing, DRM, P2P, calendaring, etc.*

A madness to the method

- *Some concepts face apparent scaling challenges*
- *The ground is too fertile – there are too many opportunities*
- *Many parts are still unproven*

The directory work

- *Incent campuses to deploy directories, and to do so in a roughly consistent fashion: The LDAP Recipe and Middleware Business Plans*
- *Create standards (syntax and semantics) for key inter-institutional attributes*
  - eduPerson
  - eduOrg
  - H.350 (nee commObj)
- *Develop tools in support of those standards*
  - LDAP Analyzer
  - Performance tools
  - Grouper

eduPerson major attributes

- *Assumes inetOrgPerson, OrgPerson*
  - Name, email address, phone, preferred language, etc.
- *Affiliations*
  - A full list (e.g. student, faculty, alum, continuing ed, etc.)
  - Primary affiliation
- *Identity (login name)*
- *Entitlements*
  - Licensed content, enrolled courses, etc.

Shibboleth

- A word which was made the criterion by which to distinguish the Ephraimites from the Gileadites. The Ephraimites, not being able to pronounce sh, called the word sibboleth. See --Judges xii.
- Hence, the criterion, test, or watchword of a party; a party cry or pet phrase.
- - Webster's Revised Unabridged Dictionary (1913):

Stage 1 - Addressing Three Scenario's

- *Member of campus community accessing licensed resource*
  - Anonymity required
- *Member of a course accessing remotely controlled resource*
  - Anonymity required
- *Member of a workgroup accessing controlled resources*
  - Controlled by unique identifiers (e.g. name)

- *Taken individually, each of these situations can be solved in a variety of straightforward ways.*
- *Taken together, they present the challenge of meeting the user's reasonable expectations for protection of their personal privacy.*

Establishing a User Context

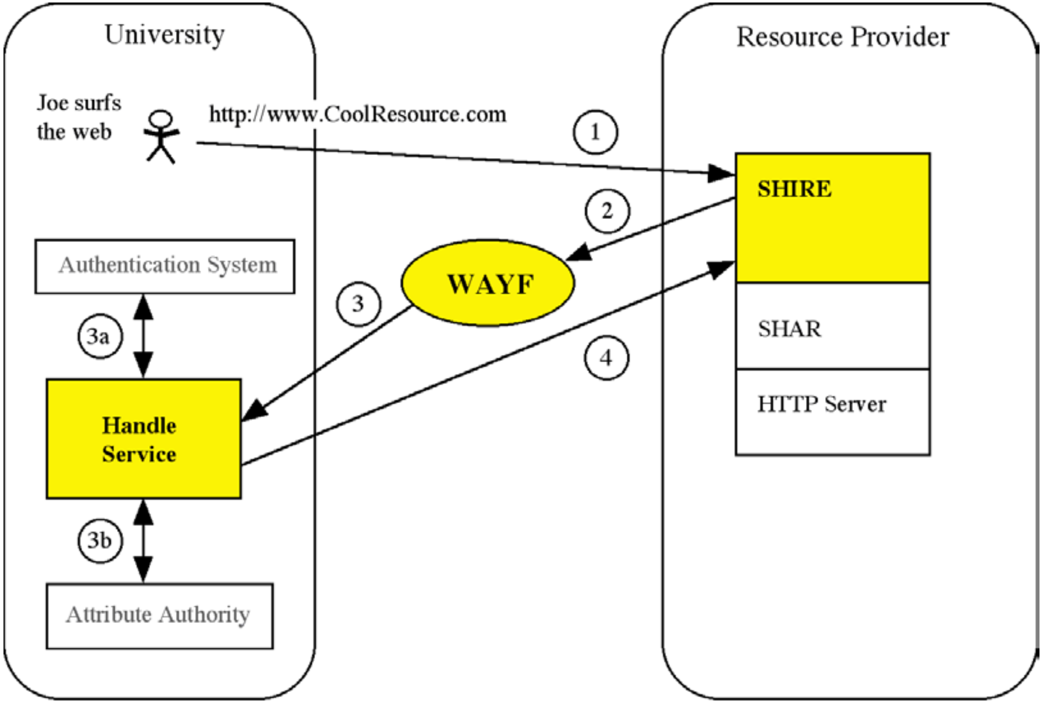

Getting Attributes and Determining Access

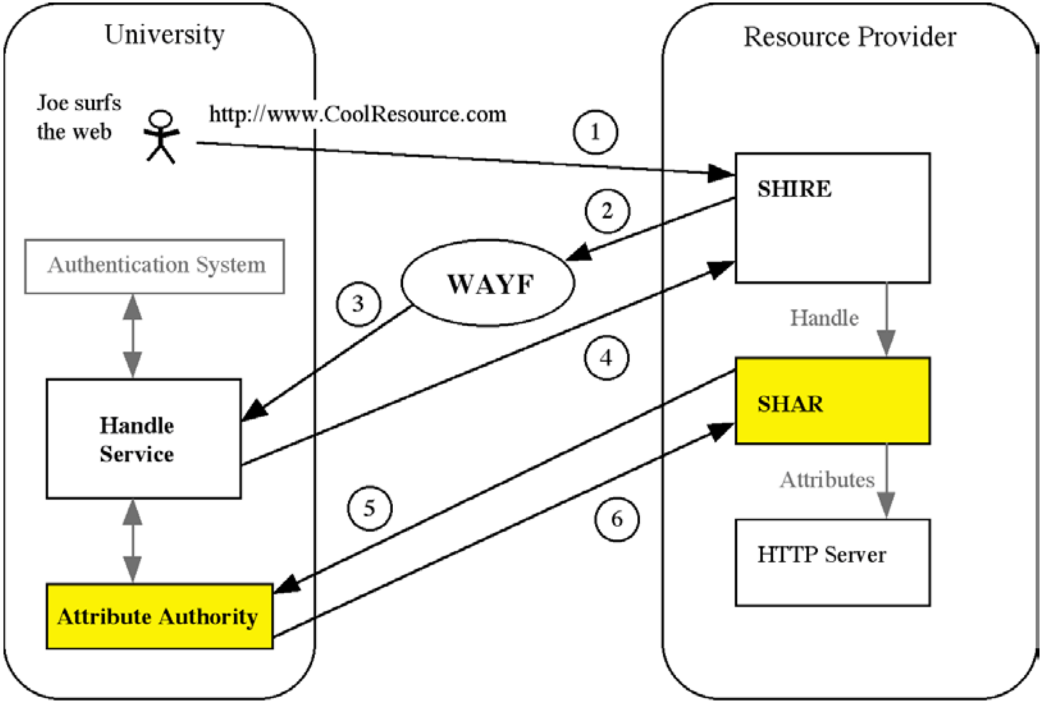

Milestones

- *Project formation - Feb 2000 Stone Soup*
- *Process - began late summer 2000 with bi-weekly calls to develop scenario, requirements and architecture.*
- *Linkages to SAML established Dec 2000*
- *Architecture and protocol completion - Aug 2001*
- *Design - Oct 2001*
- *Coding began - Nov 2001*
- *Alpha-1 release – April 24, 2002*
- *OpenSAML release – July 15, 2002*
- *v1.0 April 2003*
- *v1.1 July 2003*
- *(V2.0 early 2004)*

What are 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…*

The good

- *Very flexible – easy to establish and operate; can work for 2 or 2000 members*
- *Very customizable – tailored to fit the precise membership*
- *Address the whole problem space – security, data schema, privacy, security, transport – of inter-realm collaborations*
- *Are simple to install and operate, both for enterprises and for end-users*

The bad

- *They aren’t real, yet*
- *They don’t do everything*
- *Are web services based right now*
- *Will hit scaling walls in several dimensions; we don’t see clear answers yet…*

The unknown

- *The scaling walls*
- *How reality will unfold*
- *The convergence of the various federating software solutions*
- *Users’ willingness to manage their privacy and security*

Three Types of federation

- *Internal federations are occurring among the many subsidiaries of large companies, especially for those companies with more dynamic aggregations.*
- *Private federations occur among enterprises, typically within a market sector, that want to facilitate a specific set of transactions and interactions. Many will be bi-lateral, short-term or otherwise constrained.*
- *Public federations address more free-standing, long-term, general-purpose requirements, and need to be more open about rules of engagement. Public federations face significant scaling issues and may not be able to leverage contractual relationships that private federations can.*

Requirements for federations

- *Federation operations*
- *Federating software*
    - Exchange assertions
    - Link and unlink identities
- *Federation data schema*
- *Federation privacy and security requirements*

Federating Software

- *Liberty Alliance*
  - V 1.1 of their functional specs; 2.0 under discussion
  - Federation itself is out of scope (see PingID et al)
  - Semi-open source under development
  - Current work only on linked identities
- *Shibboleth*
  - V1.1 released; 2.0 under discussion
  - Most standards-based (though Liberty has said that they will turn their enhancements into standards organizations)
  - Pure open source
  - Current work is attribute release focused.
- *WS-**

WS-*

- *Work by Microsoft, with participation from IBM and BEA et al*
- *Complex framework, consisting of 9 areas, which can form a whole cloth solution to the problem space, but which need to closely interact with each other to do so.*
- *Several of the specifications areas still unreleased*
- *Standards process very unclear; significant IPR issues exist*
- *No implementations yet; indeed a lofty set of abstractions that will need considerable convention and detail to resolve into a working instantiation*
- *Can Shibboleth/InCommon be a working instantiation within WS-*? Good question. Once MS has all the areas defined, if someone wants to see whether the existent Shib/InCommon (or Shib/someotherfed) fits into WS-*, we'd certainly be curious…*

Shibboleth-based federations

- *InQueue*
- *InCommon*
- *Club Shib*
- *SWITCH*
- *NSDL*

- *State networks*
- *Medical networks*
- *Financial aid networks*
- *Life-long learning communities*

The Research and Education Federation Space

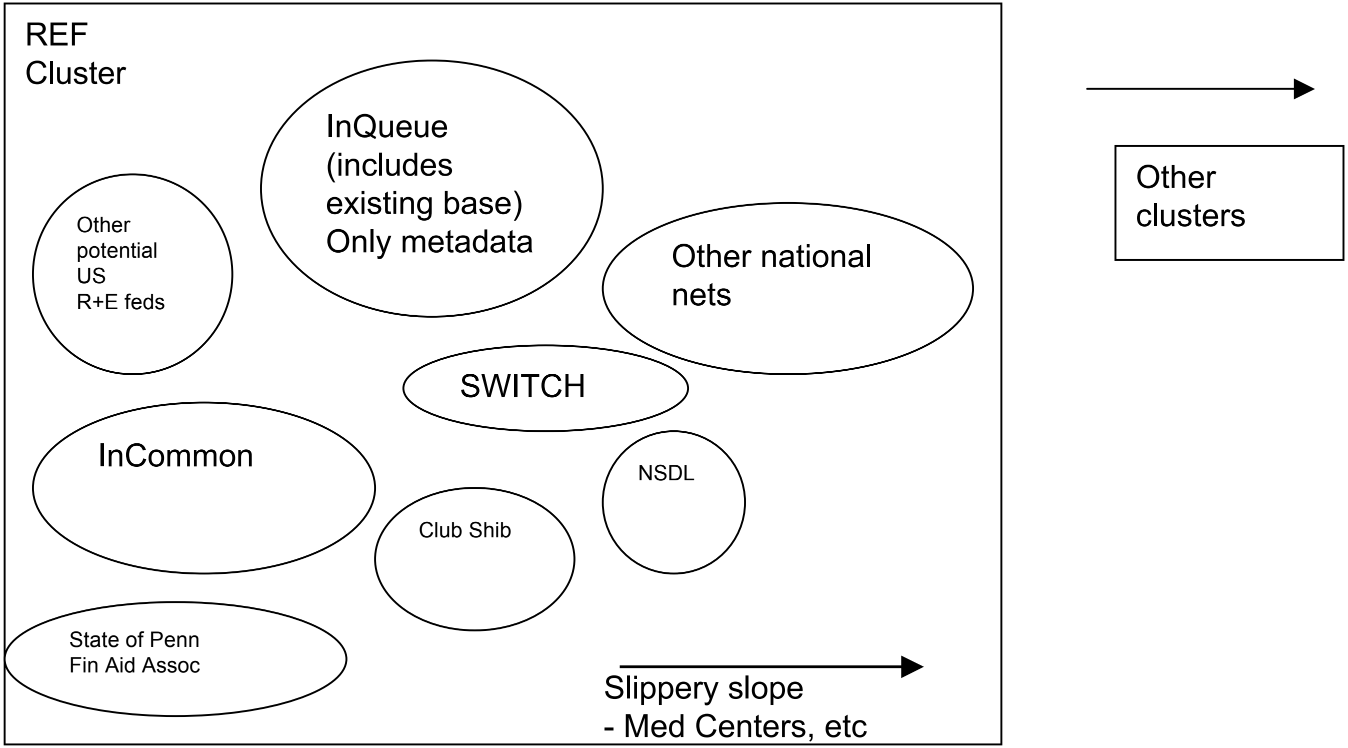

REF
Cluster

InQueue
(includes
existing base)
Only metadata

Other
potential
US
R+E feds

Other national
nets

SWITCH

InCommon

NSDL

Club Shib

State of Penn
Fin Aid Assoc

Slippery slope
- Med Centers, etc

Other
clusters

Trust pivot points in federations

- *In response to real business drivers and feasible technologies*

  increase the strengths of

  - Campus/enterprise identification, authentication practices
  - Federation operations, auditing thereof
  - Campus middleware infrastructure in support of Shib (including directories, attribute authorities and other Shib components) and auditing thereof
  - Relying party middleware infrastructure in support of Shib
  - Moving in general from self-certification to external certification

Collaboration Tools

- *Personal Privacy and Resource Managers*
- *Digital rights management*
- *Role-based access controls*
- *Desktop videoconferencing*
- *Interrealm calendaring*
- *Authenticated instant messaging*
- *P2P*
- *Shibbed **

Personal Resource Manager

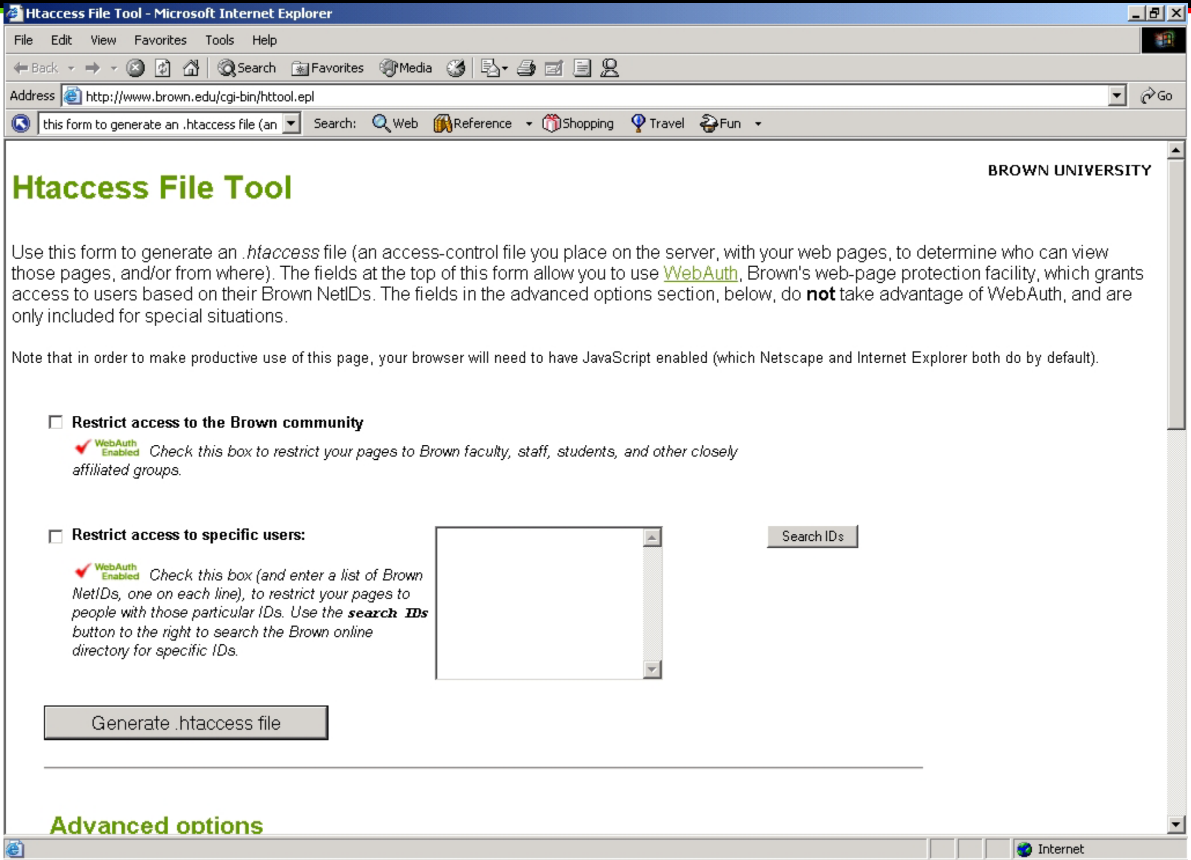

Htaccess File Tool - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://www.brown.edu/cgi-bin/httool.epl
Go
this form to generate an .htaccess file (an
Search: Web Reference Shopping Travel Fun
Internet

BROWN UNIVERSITY

Htaccess File Tool

Use this form to generate an *.htaccess* file (an access-control file you place on the server, with your web pages, to determine who can view those pages, and/or from where). The fields at the top of this form allow you to use WebAuth, Brown's web-page protection facility, which grants access to users based on their Brown NetIDs. The fields in the advanced options section, below, do **not** take advantage of WebAuth, and are only included for special situations.

Note that in order to make productive use of this page, your browser will need to have JavaScript enabled (which Netscape and Internet Explorer both do by default).

□ **Restrict access to the Brown community**

✓ WebAuth Enabled *Check this box to restrict your pages to Brown faculty, staff, students, and other closely affiliated groups.*

□ **Restrict access to specific users:**

✓ WebAuth Enabled *Check this box (and enter a list of Brown NetIDs, one on each line), to restrict your pages to people with those particular IDs. Use the **search IDs** button to the right to search the Brown online directory for specific IDs.*

Search IDs

Generate .htaccess file

Advanced options

Privacy Management Systems

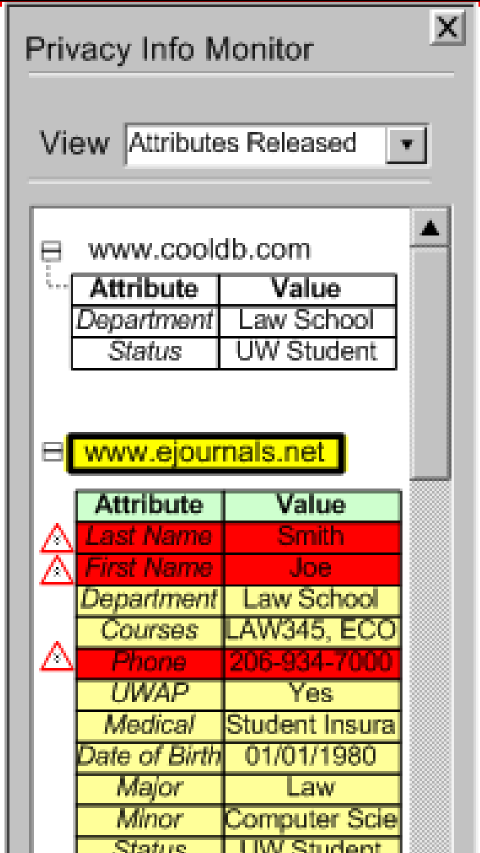

View: Attributes Released

www.cooldb.com

Attribute	Value
Department	Law School
Status	UW Student

www.ejournals.net

	Attribute	Value
⚠	*Last Name*	Smith
⚠	*First Name*	Joe
	Department	Law School
	Courses	LAW345, ECO
⚠	*Phone*	206-934-7000
	UWAP	Yes
	Medical	Student Insura
	Date of Birth	01/01/1980
	Major	Law
	Minor	Computer Scie
	Status	UW Student

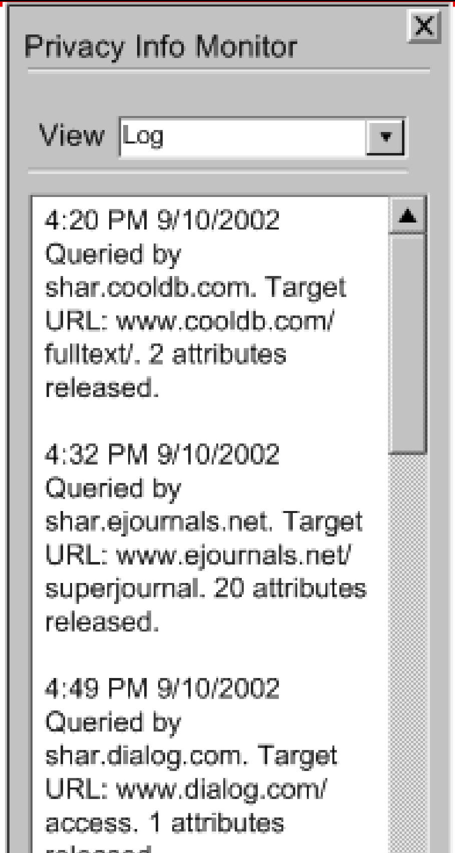

View: Log

4:20 PM 9/10/2002 Queried by shar.cooldb.com. Target URL: www.cooldb.com/fulltext/. 2 attributes released.

4:32 PM 9/10/2002 Queried by shar.ejournals.net. Target URL: www.ejournals.net/superjournal. 20 attributes released.

4:49 PM 9/10/2002 Queried by shar.dialog.com. Target URL: www.dialog.com/access. 1 attributes released.

PMS-2

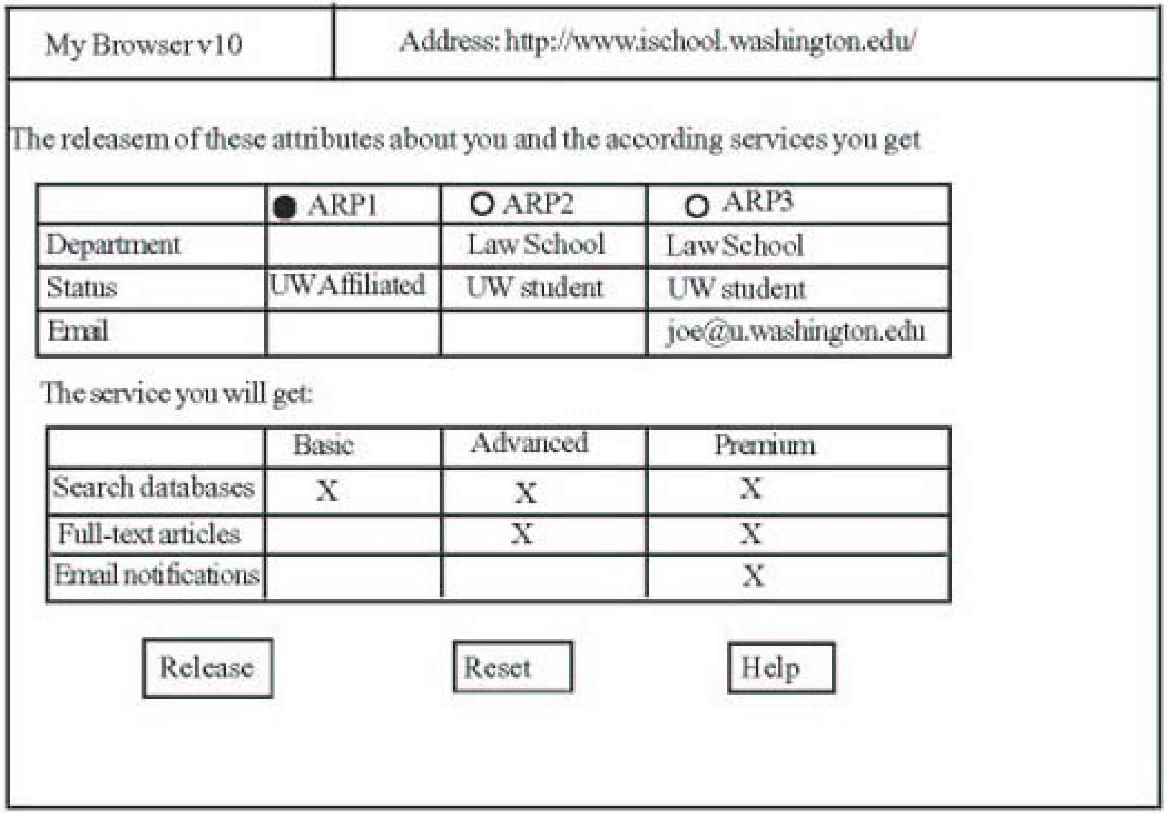

My Browser v10

Address: http://www.ischool.washington.edu/

The releasem of these attributes about you and the according services you get

	● ARP1	○ ARP2	○ ARP3
Department		Law School	Law School
Status	UW Affiliated	UW student	UW student
Email			joe@u.washington.edu

The service you will get:

	Basic	Advanced	Premium
Search databases	X	X	X
Full-text articles		X	X
Email notifications			X

Access to Digital Materials

- *Several ways to use digital materials –*
- *personal use – typically purchased by individuals on a subscription or per-use basis.*
- *professional use – typically acquired (for fee or legal agreement) by an organization or university on a bulk basis, with access redistributed freely to members of the organization.*
- *public use – as a citizen, entitled to an information commons, and other basic information rights, such as Fair Use and Freedom of Information*

Digital rights technologies

- *The different uses of on-line materials have different requirements; they will likely require different technologies.*

- *Requirements vary about the needs and controls for privacy, the economic recovery model, the needs and controls for security, etc.*

- *Who is developing the digital rights technologies for professional and public use?*

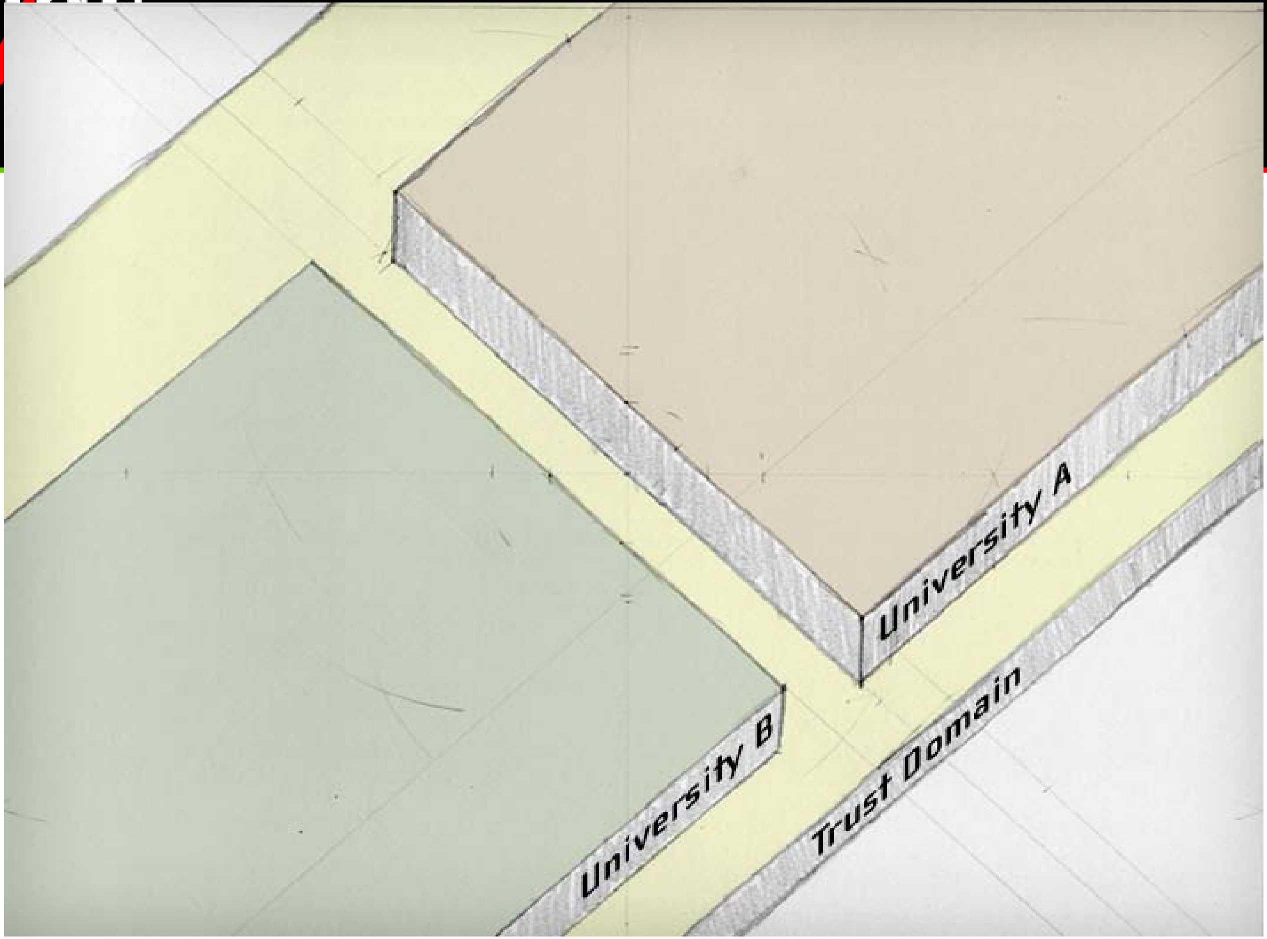
University A
University B
Trust Domain

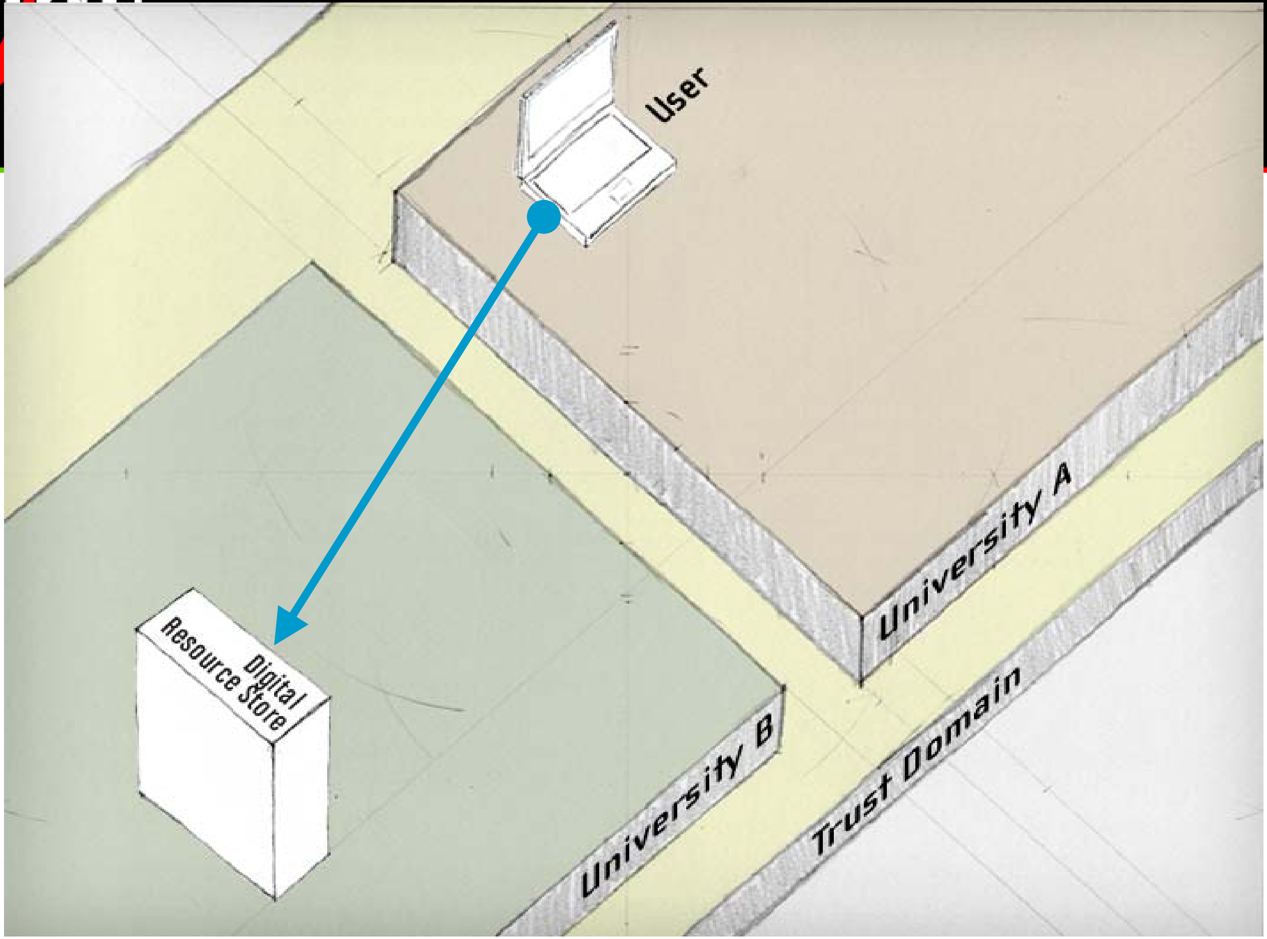
User
University A
Digital Resource Store
University B
Trust Domain

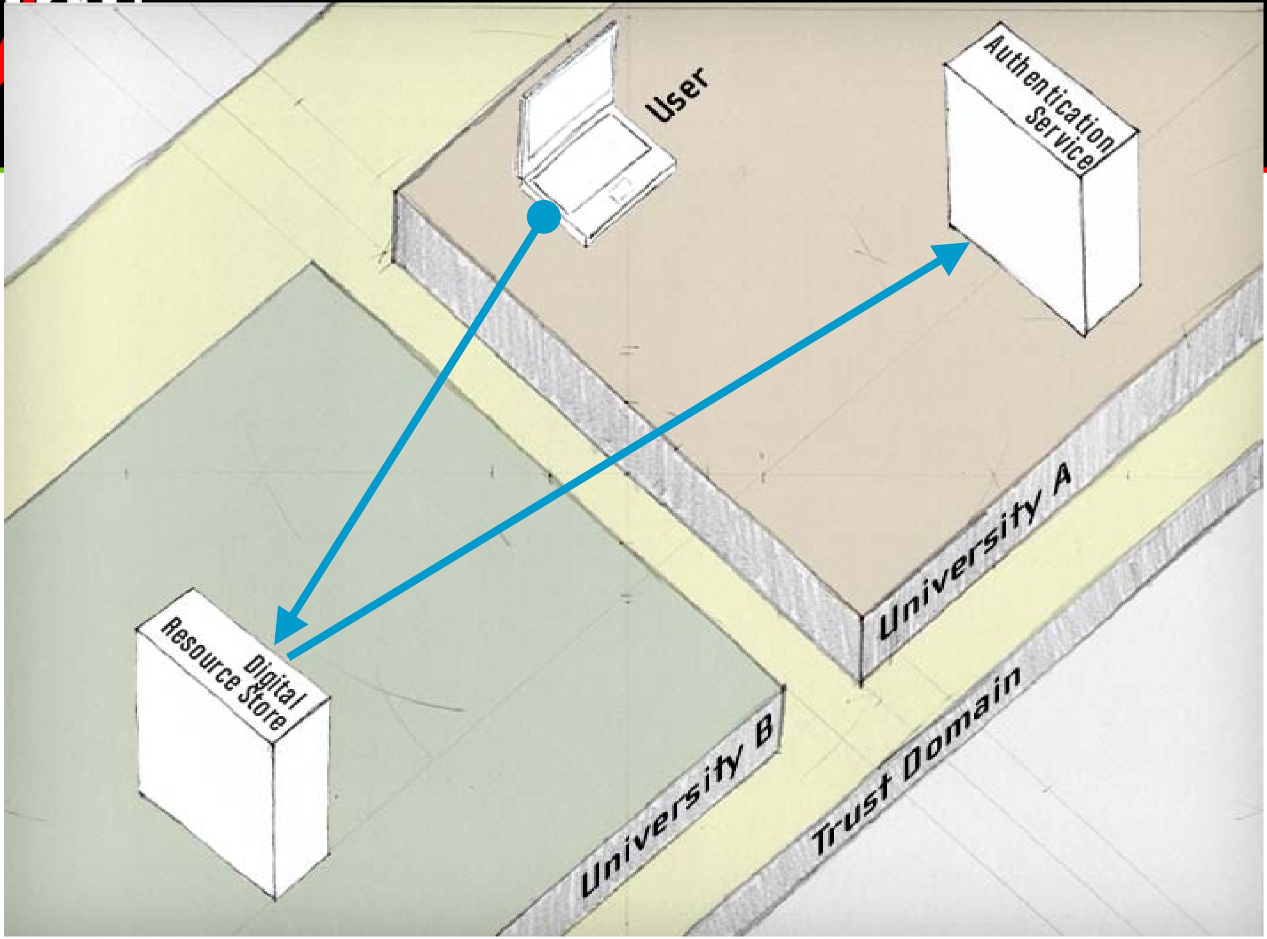
User
Authentication Service
Digital Resource Store
University A
University B
Trust Domain

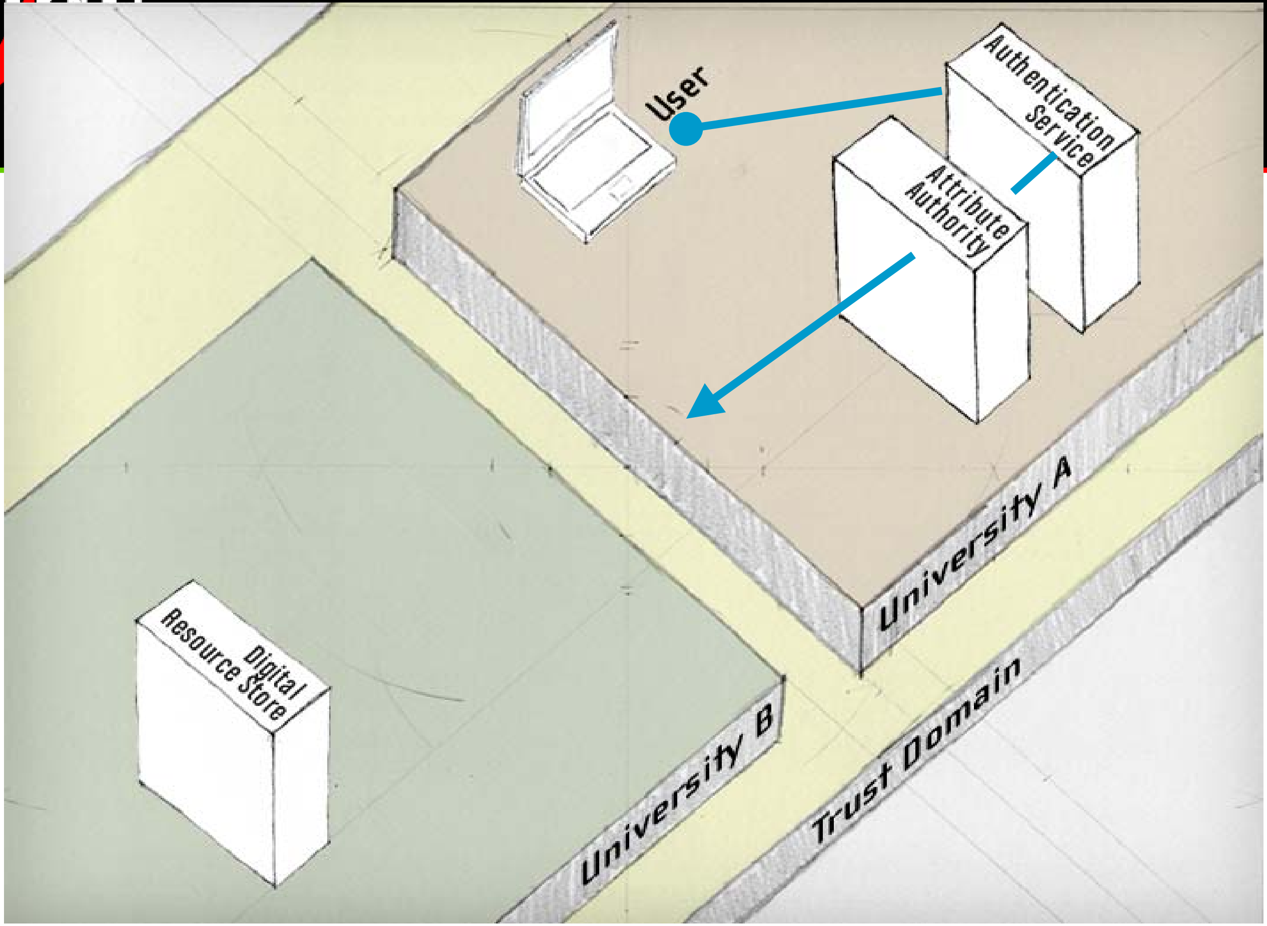
User
Authentication Service
Attribute Authority
Digital Resource Store
University A
University B
Trust Domain

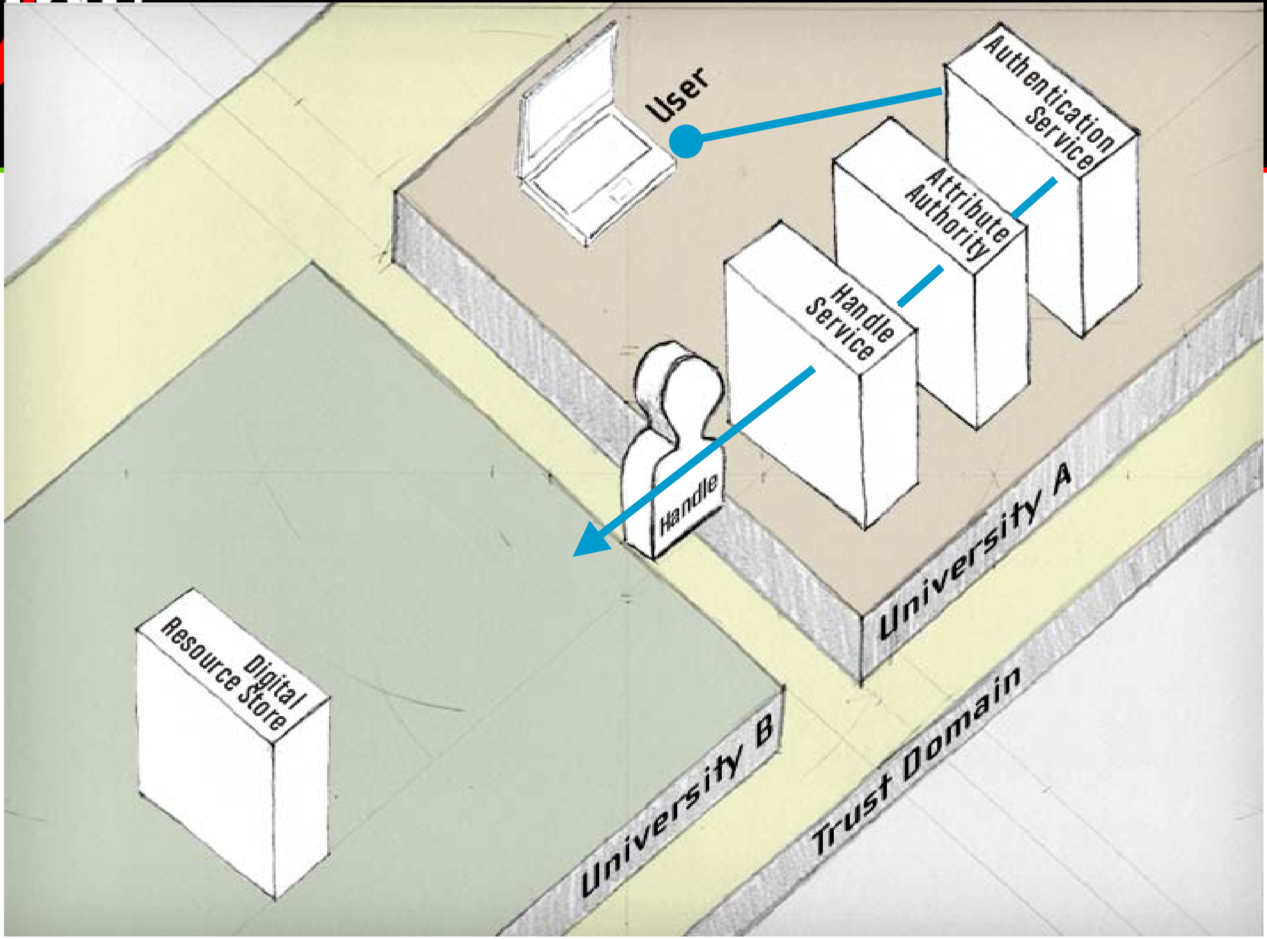
User
Authentication Service
Attribute Authority
Handle Service
Handle
University A
Digital Resource Store
University B
Trust Domain

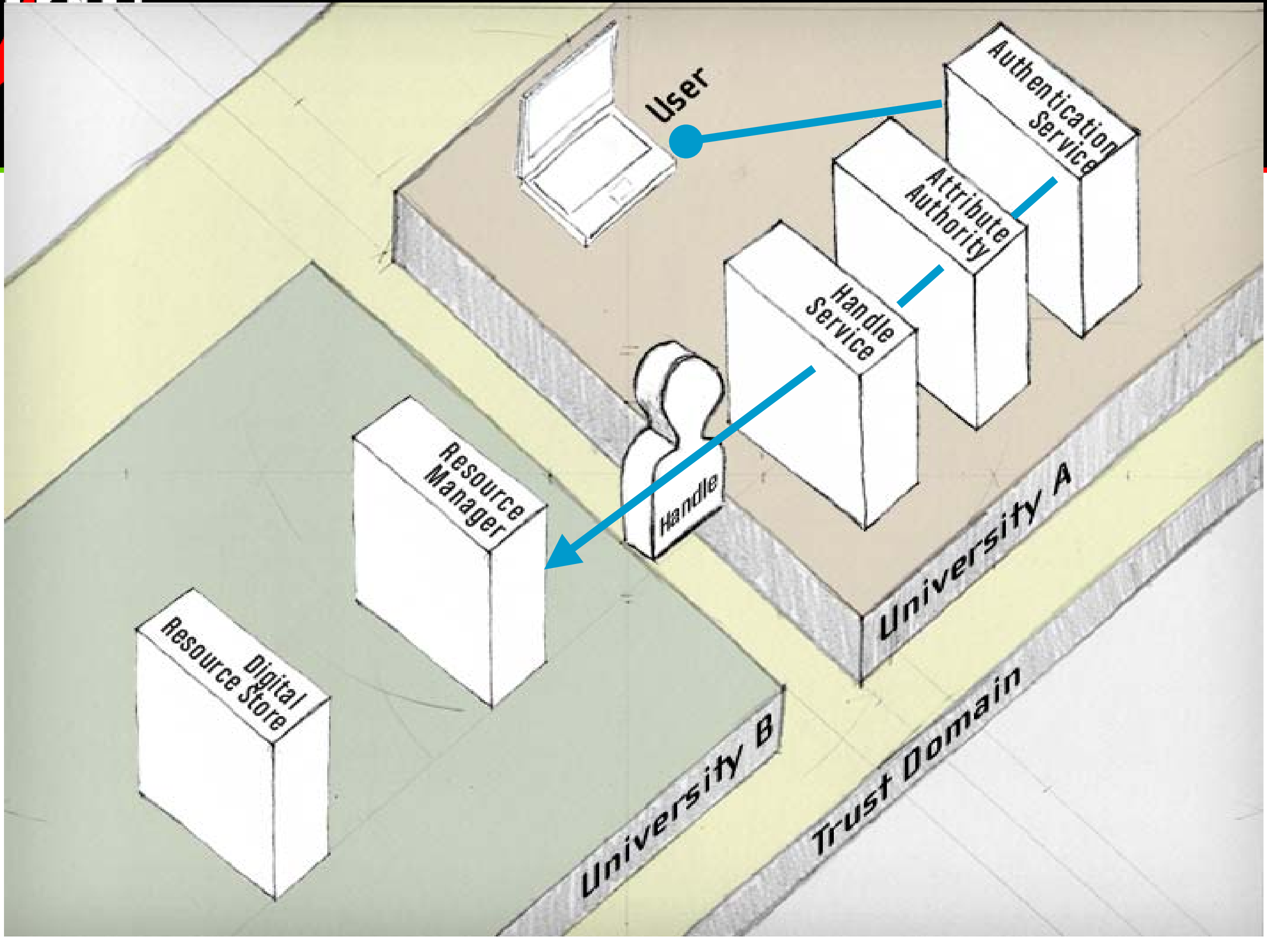
User
Authentication Service
Attribute Authority
Handle Service
Handle
Resource Manager
Digital Resource Store
University A
University B
Trust Domain

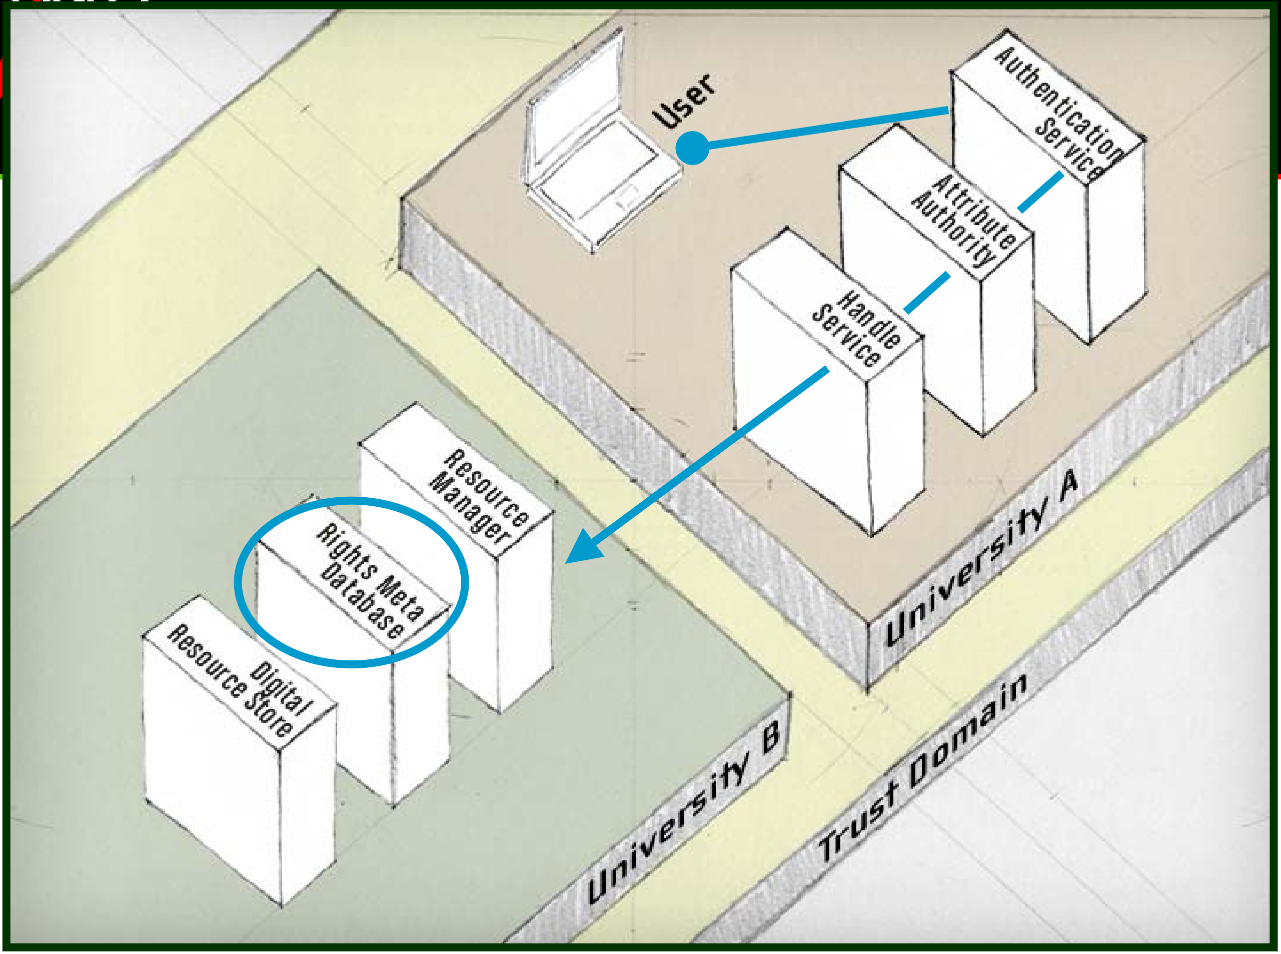
User
Authentication Service
Attribute Authority
Handle Service
Resource Manager
Rights Meta Database
Digital Resource Store
University A
University B
Trust Domain

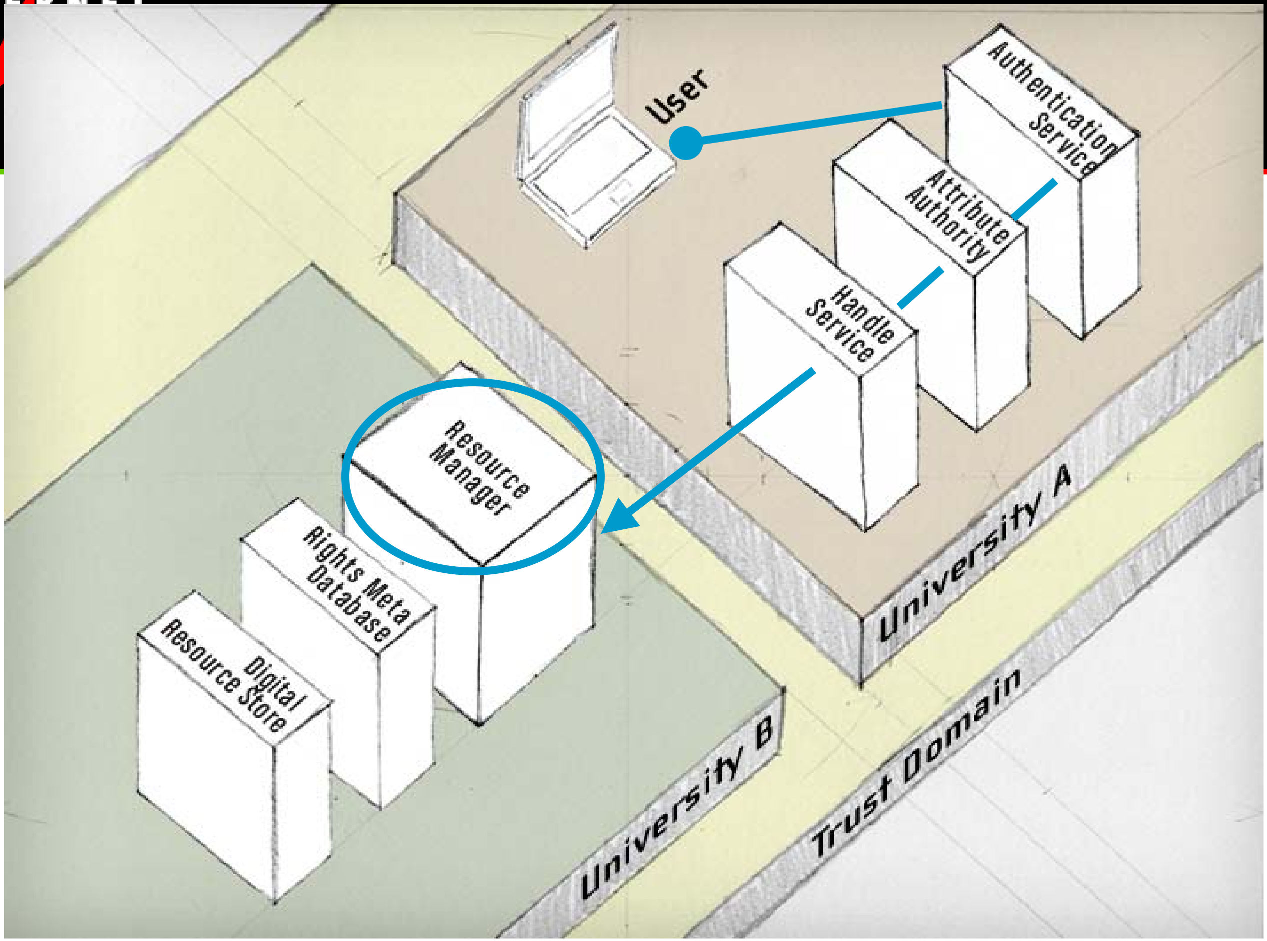
User
Authentication Service
Attribute Authority
Handle Service
University A
Resource Manager
Rights Meta Database
Digital Resource Store
University B
Trust Domain

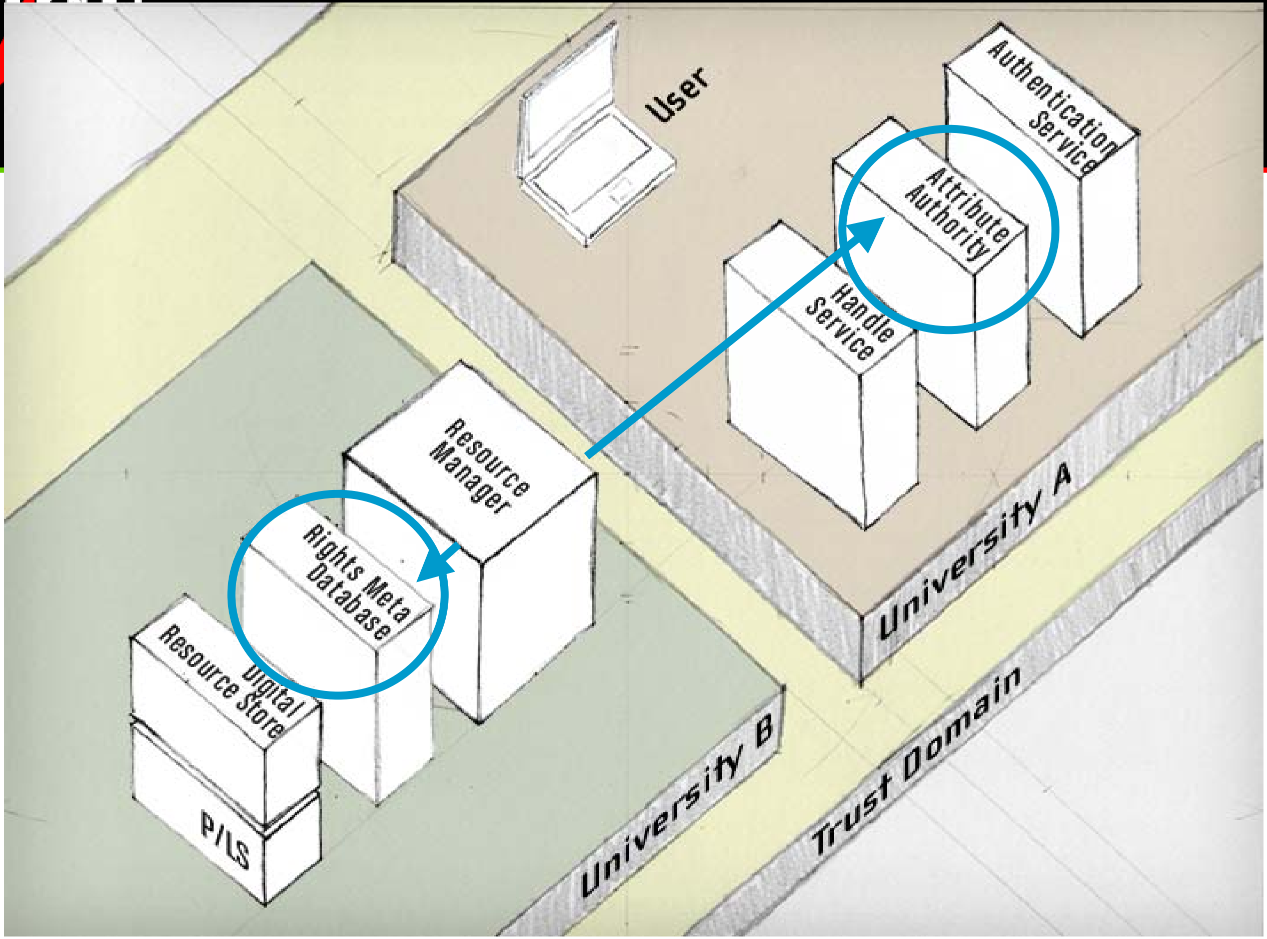
User
Authentication Service
Attribute Authority
Handle Service
Resource Manager
Rights Meta Database
Digital Resource Store
P/LS
University A
University B
Trust Domain

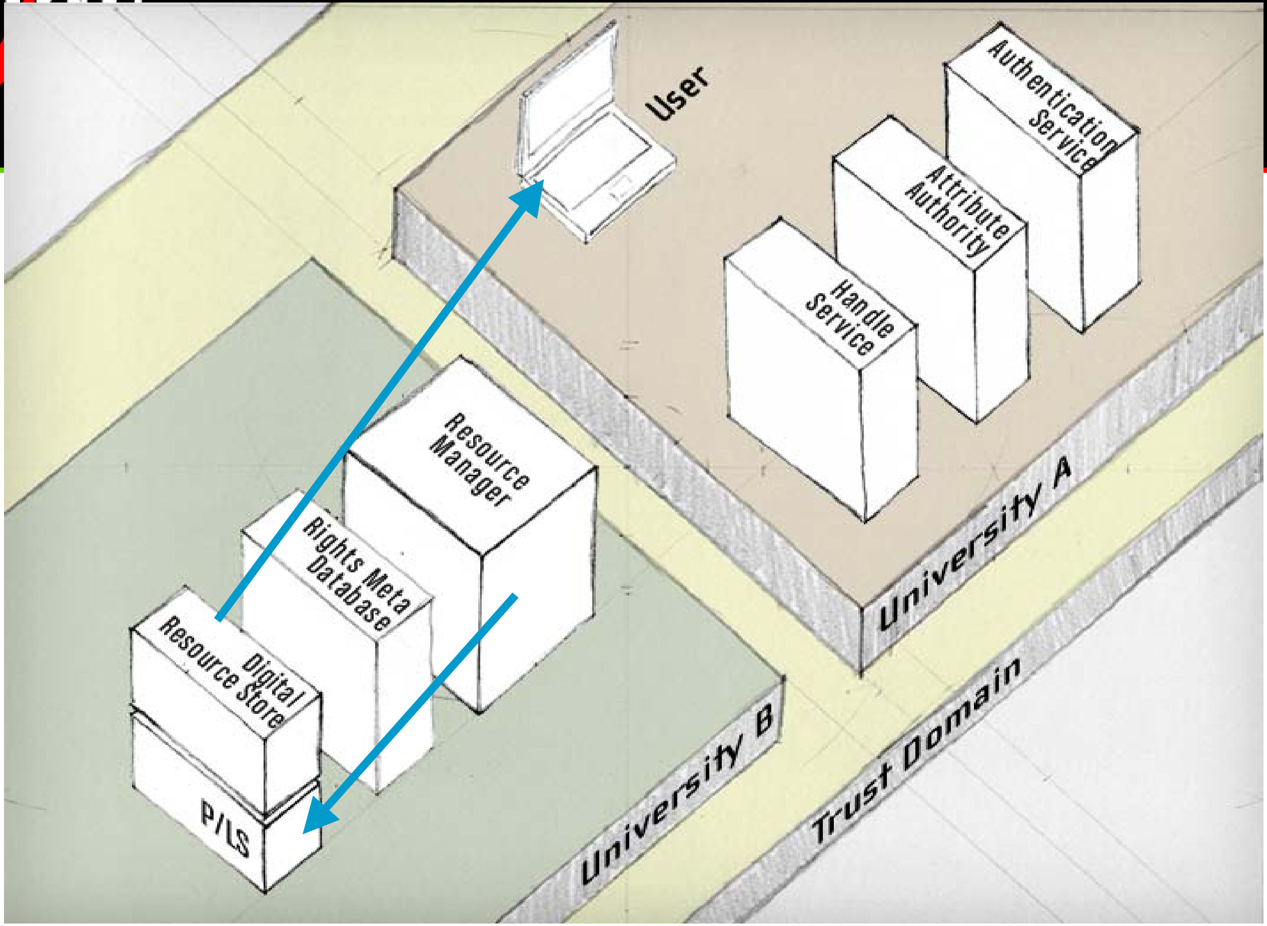
User
Authentication Service
Attribute Authority
Handle Service
University A
Resource Manager
Rights Meta Database
Digital Resource Store
P/LS
University B
Trust Domain

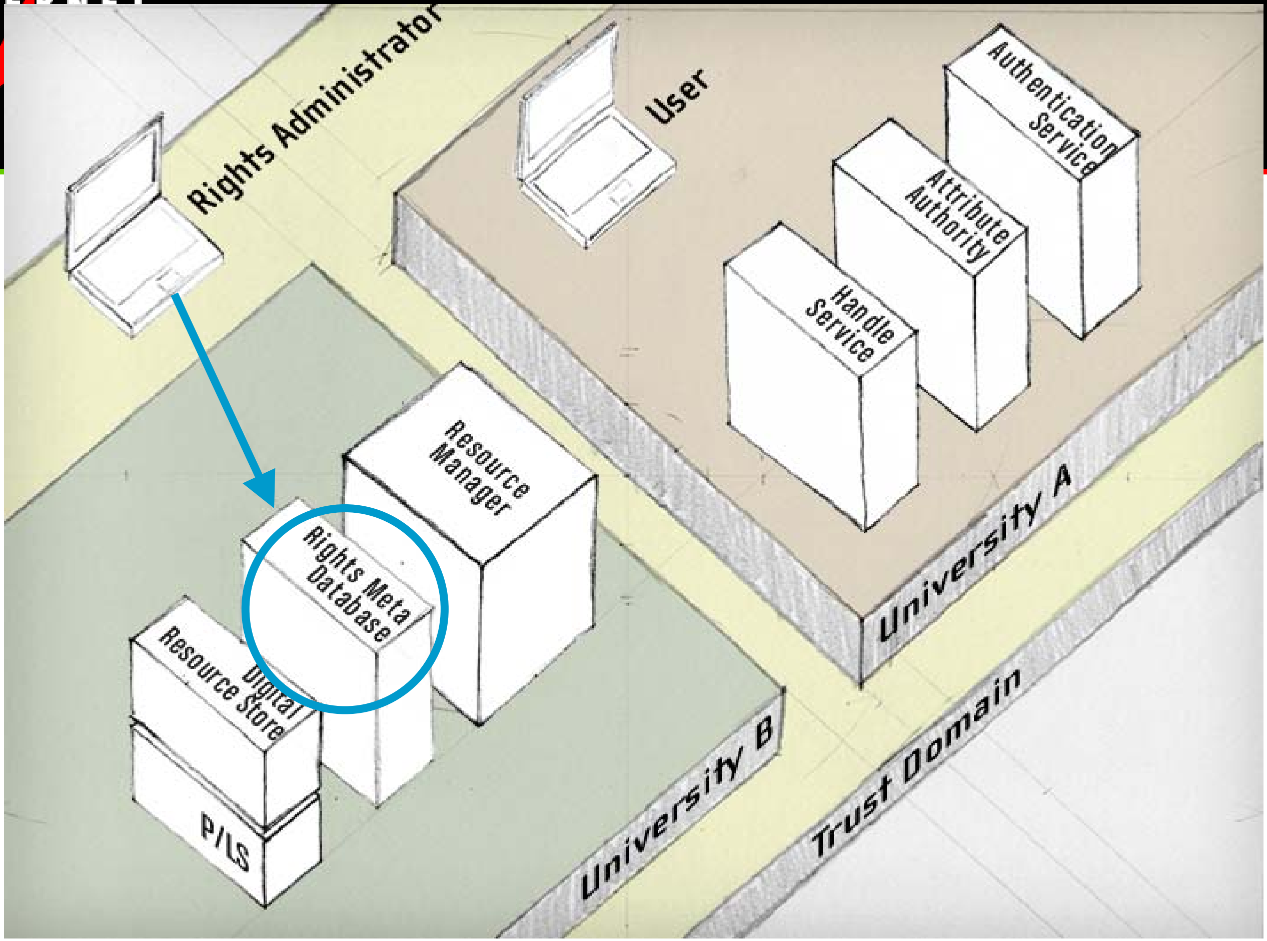
Rights Administrator
User
Authentication Service
Attribute Authority
Handle Service
Resource Manager
Rights Meta Database
Digital Resource Store
P/LS
University A
University B
Trust Domain

Simpleton's Authz

- *Mt Authorization, from whose top one sees the unified field theory of authzanity*

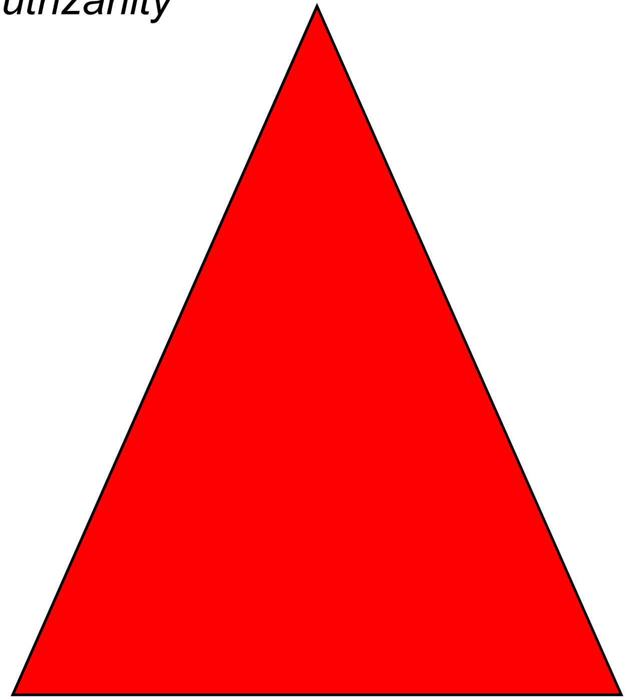

A Better View of Authzanity

- *The Enterprise Attribute River,*
- *feeding the legacy apps that plumb to it, for their own internal authz processing*
- *The Oracle in the Mountains*
- *providing authz advice/decision (permit, deny, permit with constraints, etc.) to the simple masses*
- *The Policy Decision Point*
- *The Policy Enforcement Points*
- *throughout the land, implementing decisions*
- *The Interrealm Attribute Requestor*
- *The Packaging Plants*
- *for digital rights execution, etc.*

The Enterprise Attribute River

- *Stanford University Authority Project*
- *http://www.stanford.edu/group/itss-ccs/project/authority/*
- *Key individuals: Lynn McRae, Sandy Senti, the lingering vapors of Bob Morgan and Jeff Hodges…*
- *Marshalls resources from a broad set of registries (for people, groops, policies, roles, etc) and directories, with innovative use of groups, to produce atomic entitlements for entry into existing authorization systems*
- *Provides users with facile tools for viewing, managing and delegating authority*
- *Requires reasonable alignment of roles within institution.*
-

Stanford Authz Model

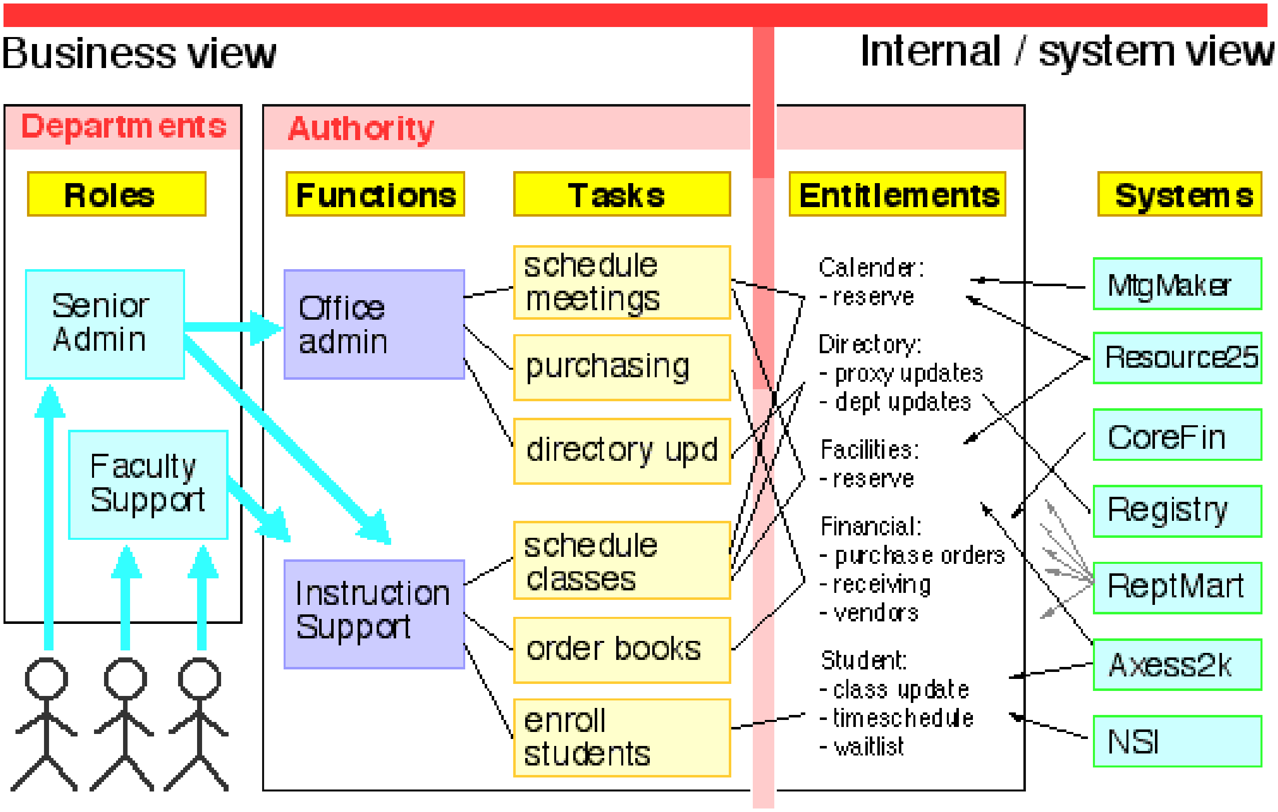

Stanford Authz Goals

- ***Simplification** of authority policy, management and interpretation. We should be able to summarize the full rights and privileges of an individual "at a glance" or let departmental administrators view and manage together all authority in their department or division.*
- ***Consistent** application of authority rules via infrastructure services and synchronization of administrative authority data across systems.*
- ***Integration** of authority data with enterprise reference data to provide extended services such as delegation and automatic revocation of authority based on status and affiliation changes.*
- ***Role-based** authority, that is, management of privileges based on job function and assignments rather than attached to individuals.*

Deliverables

- *The deliverables consist of*
- *a Web interface and program APIs to provide distributed management (to the departments, to external programs) of access rights and privileges, and*
- *delivery of authority information through the infrastructure as directory data and authority events.*

Desktop video

- *Point-and-click initiation*
  - Directories – H.350
  - Presence
- *Authentication*
  - For Origin Control
  - For Target Control
- *Authorization*
- *Across H.323, SIP, AG, VRVS,…*

VidMid-VC Scenario [1]

- Origin user initiates VC session by selecting from:
  - Local directory (“White Pages”), or
  - Federated directory(s), or
  - Drop-down list of frequent targets
    - Directory-enabled -- always up to date
- Target user entry denotes client capability(s):
  - H.323
  - SIP
  - Voice

VidMid-VC Scenario [1a]

- GUI tools to enable user to manage their location to receive VC, e.g.:
  - For the next 90 minutes I will be in Conf. Room B, if anyone tries to reach me
  - For the next 4 hours I will be at 192.168.1.244, if Jane Doe, Bob Smith, or Alice Jones tries to reach me. All others should be notified that I will be available for VC at my own station at 14:00 today.
  - I want calls from Prof. Jones from the Medical School to be encrypted

VidMid-VC Scenario [2]

- Authentication

- Utilize user's authentication credential from their "home" security domain
- On the target side, pop-up window appears on workstation, reading:
  - Jane Doe from Penn State U. is attempting to initiate a videoconference with you, and Penn State U. asserts that this is in fact Jane Doe. Would you like to accept the call? [Yes] [No]
- Needed: ability to utilize credentials received from local security domain, and pass them as needed to Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs) as needed

VidMid-VC Scenario [3]

- Authorization
- Management tools to control access to VC based on role and/or location. E.g.:
  - Dormitory student users not permitted to VC between 08:00–17:00 Mon-Fri
  - Math Dept. Admin Asst. permitted to initiate VC on behalf of faculty/staff at any time.
  - Only faculty/staff users permitted to exceed 384 Kbps
  - Faculty is allowed to set up a multipoint conference call
  - GUI tools to enable user to manage their location to receive VC
- Need clients enabled to be controlled based on user role and/or device location, or other attributes

Issues

- *Virtual Organizations and Federations and the OCLC community*
- *Role-Based Access Control Opportunities*
- *Entitlements and attributes: target and origin PDP’s*
- *The Unified Field Theory of Trust and Its Meaning for Metadata*

Issues and turtles

- *The Bertrand Russell turtle*
- *will the layers get to be too much*
- *will the complexity be manageable by users*
- *The Edward Oppenheimer turtle*
- *anticipate unintended consequences*