

Access Information Management

Tom Barton

University of Chicago



Topics

- Shibboleth & Federations
- Federations & Grids: the Long Tail
- Federated Identity & TeraGrid
- Managing access for VO-like things
 - COmanage, GridGrouper

Shibboleth Status

- Shib 1.3 the widely deployed base
- OpenSAML 2.0 libraries broadly used
- Shib 2.0 now in beta
- “Shib 2.0 will interoperate with other SAML 2.0 products better than they interoperate with each other.”
- NSF, Internet2, JISC, SWITCH, Google and MS, among others have provided funding
- Support services businesses developing in the US and overseas

Shibboleth use

- ~12 M in Europe/Asia and ~6 M in the US; growing exponentially in many countries; almost all Shib 1.3
- Almost all users do not know they are using it (some may see a redirect...) but that is to change
- InCommon, Texas (three federations), UCTrust, CalStateTrust, CCLA of Florida, CC of Washington State
- DHS + DOJ, Dept of Ed
- OpenSAML used by Google, Verisign, etc.

The rise of federations

- Federations are now occurring broadly, and internationally, to support inter-institutional and external partner collaborations
- Almost all in the corporate world are bilateral; almost all in the R&E world are multilateral
- They provide a powerful leverage of enterprise (campus, site) credentials
- Federations are learning to peer
- Internal federations are also proving quite useful

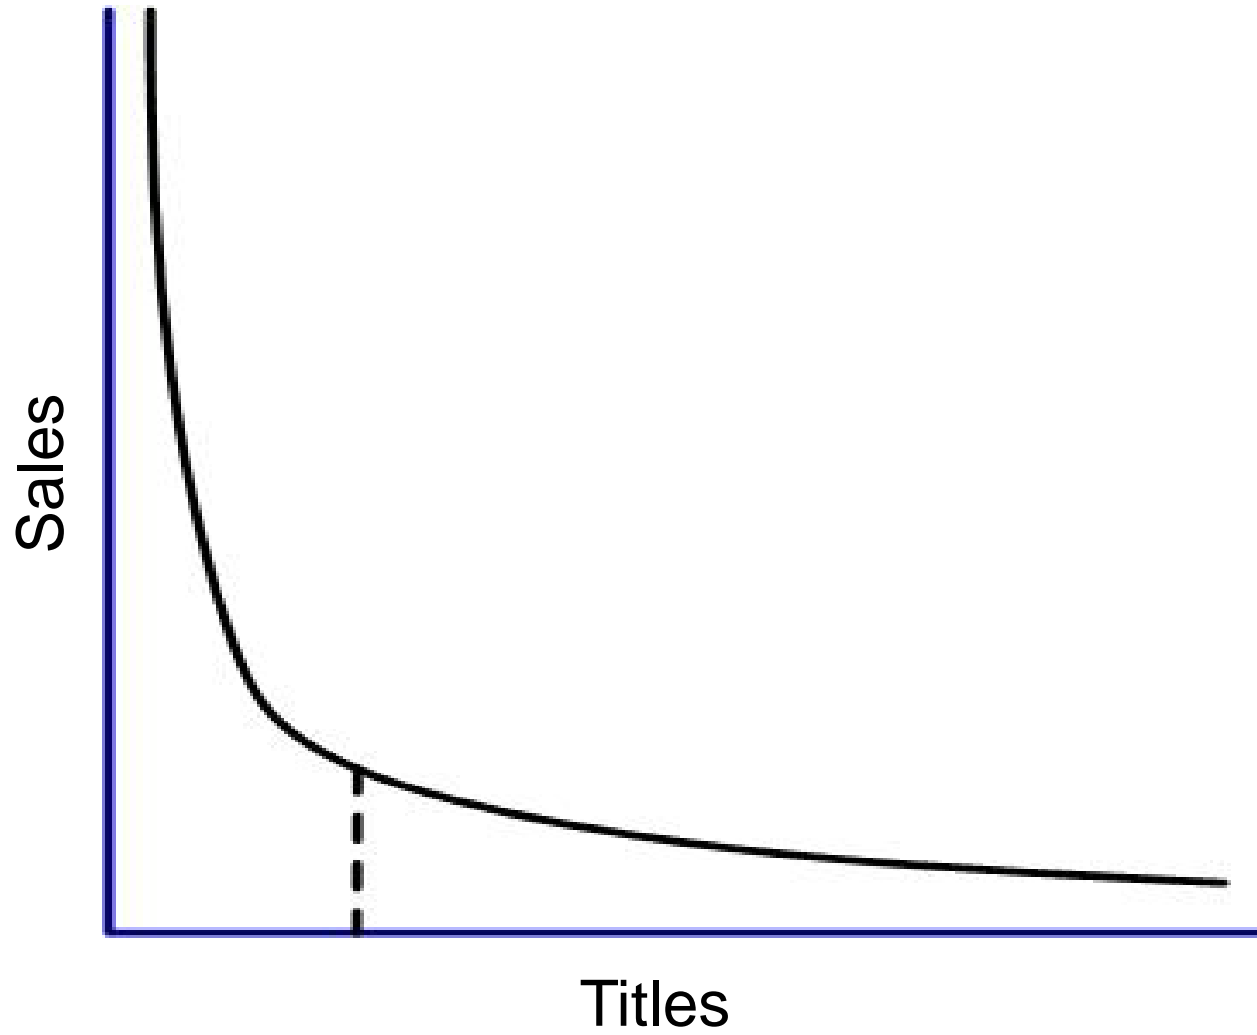
InCommon Federation: Essential Data

- US R&E Federation, a 501(c)3
- Addresses legal, LOA, shared attributes, business proposition, etc issues
- Members are universities, service providers, government agencies, national labs
- Over 70 organizations and growing steadily; 1.3 million user base now, crossing 2 million by the end of the year
- Use ranges over popular and academic content, wiki and list controls, ASPs, NIH applications, ...
- www.incommonfederation.org

Prague Meeting on Inter-federation

- 15-20 international R&E federations (5 continents) plus Liberty Alliance and a few others
- Prague, September 3
- Lots of topics: Attribute mapping, Privacy Policies, Dispute resolution, Financial considerations, Technical direction setting
- Next steps:
 - UK drafting an analysis of International Peering needs, opportunities, etc.
 - Discussions with Liberty EGovSIG (e.g. SAML 2.0 profiles, attribute schema)

The Long Tail

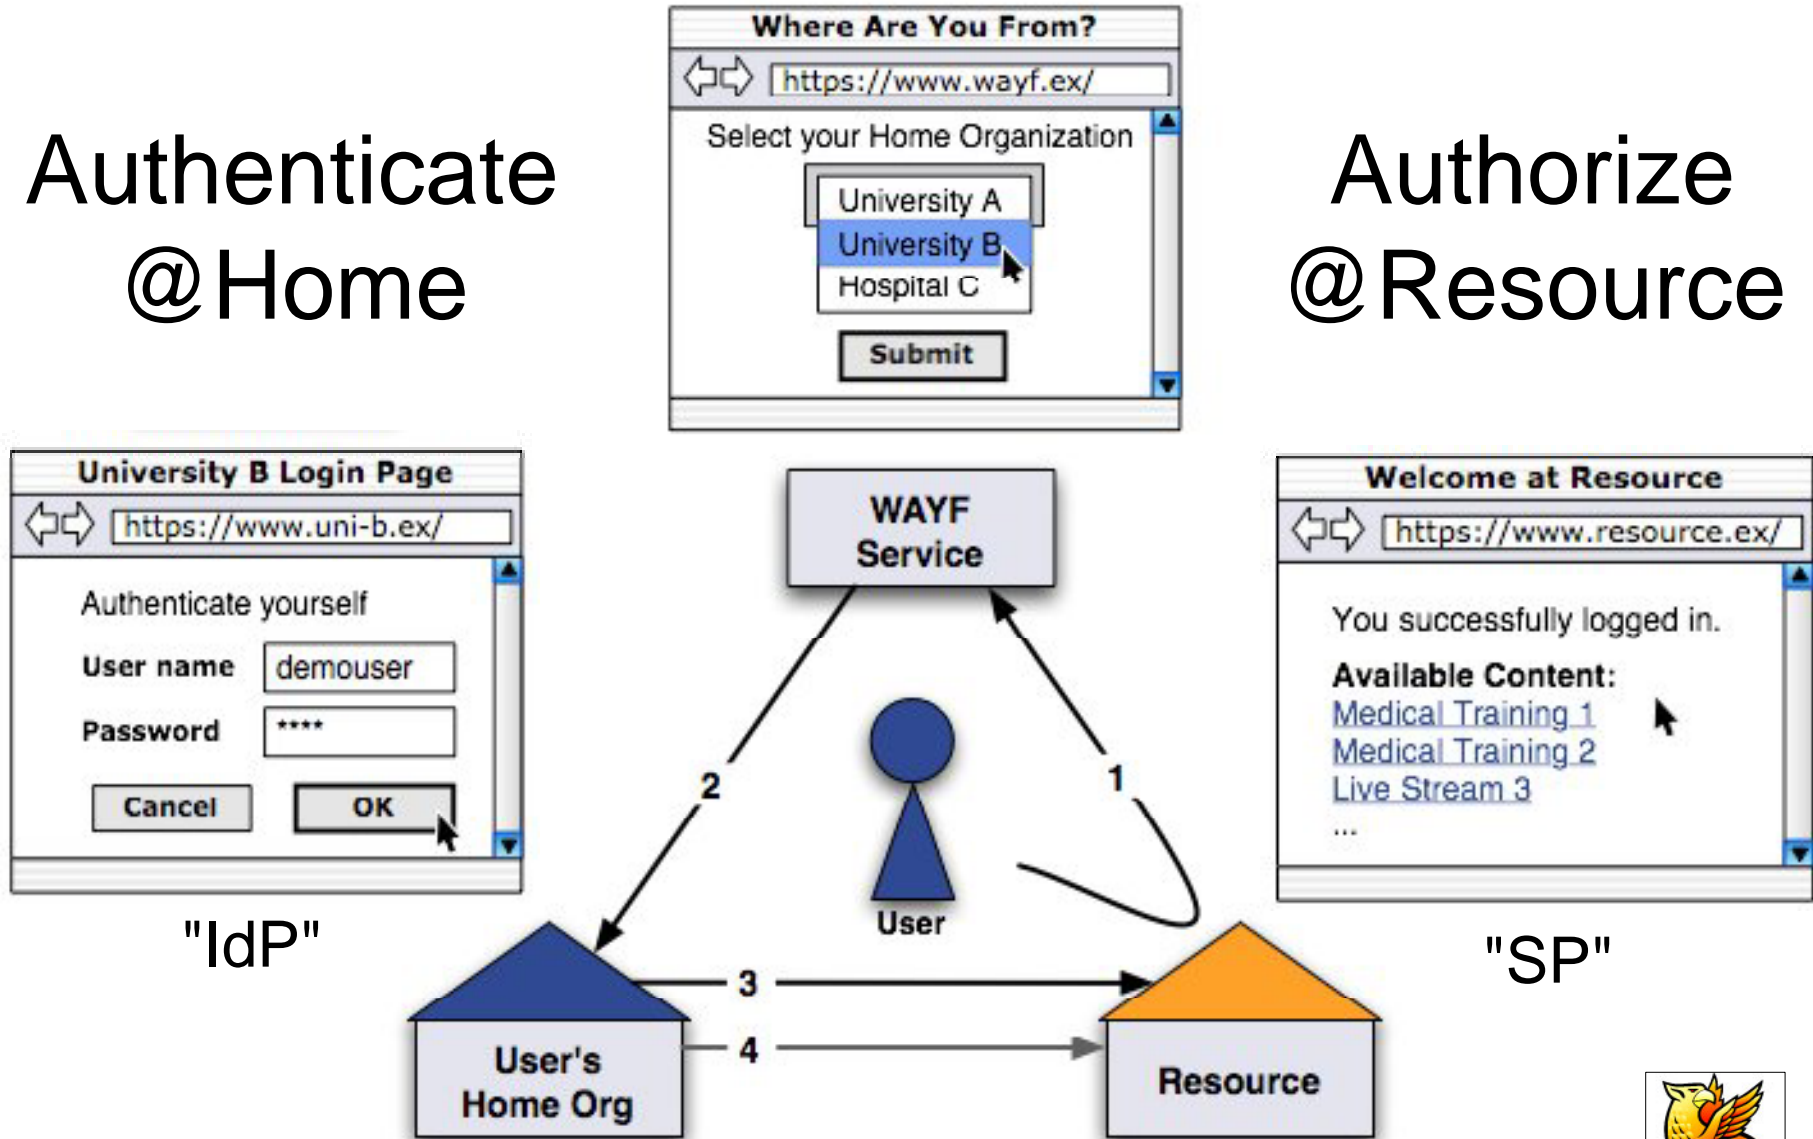


Scaling TeraGrid Usership



Authenticate @Home

Authorize @Resource



Federated Identity



ala Shibboleth

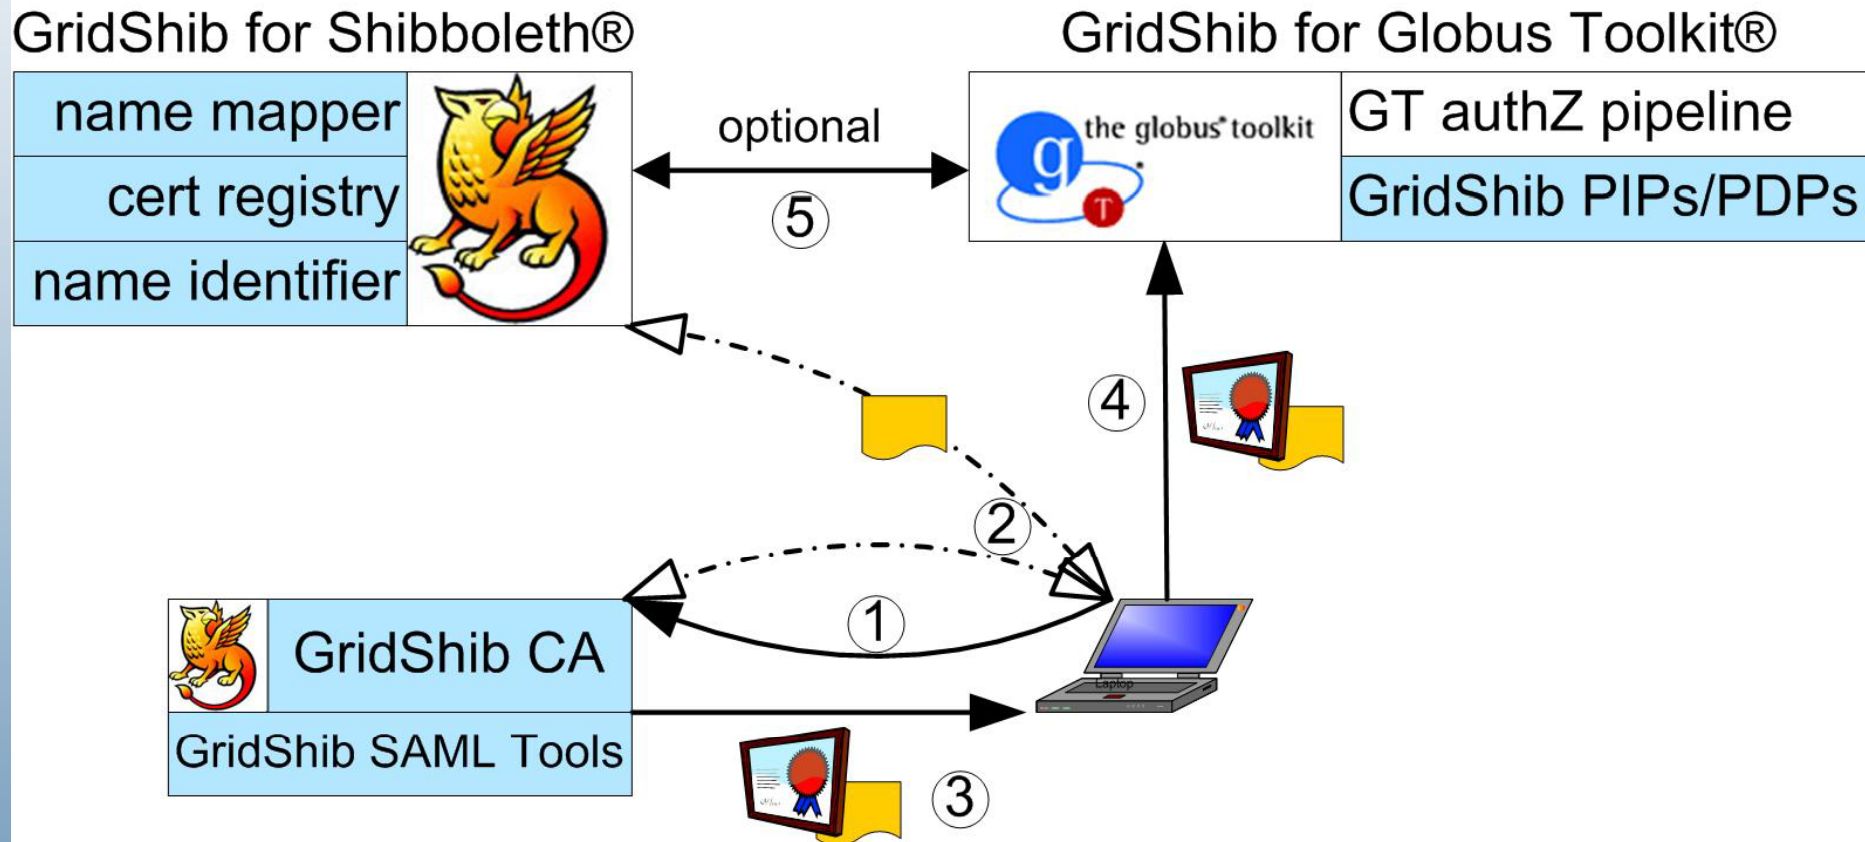
InCommon Federation: Essential Services

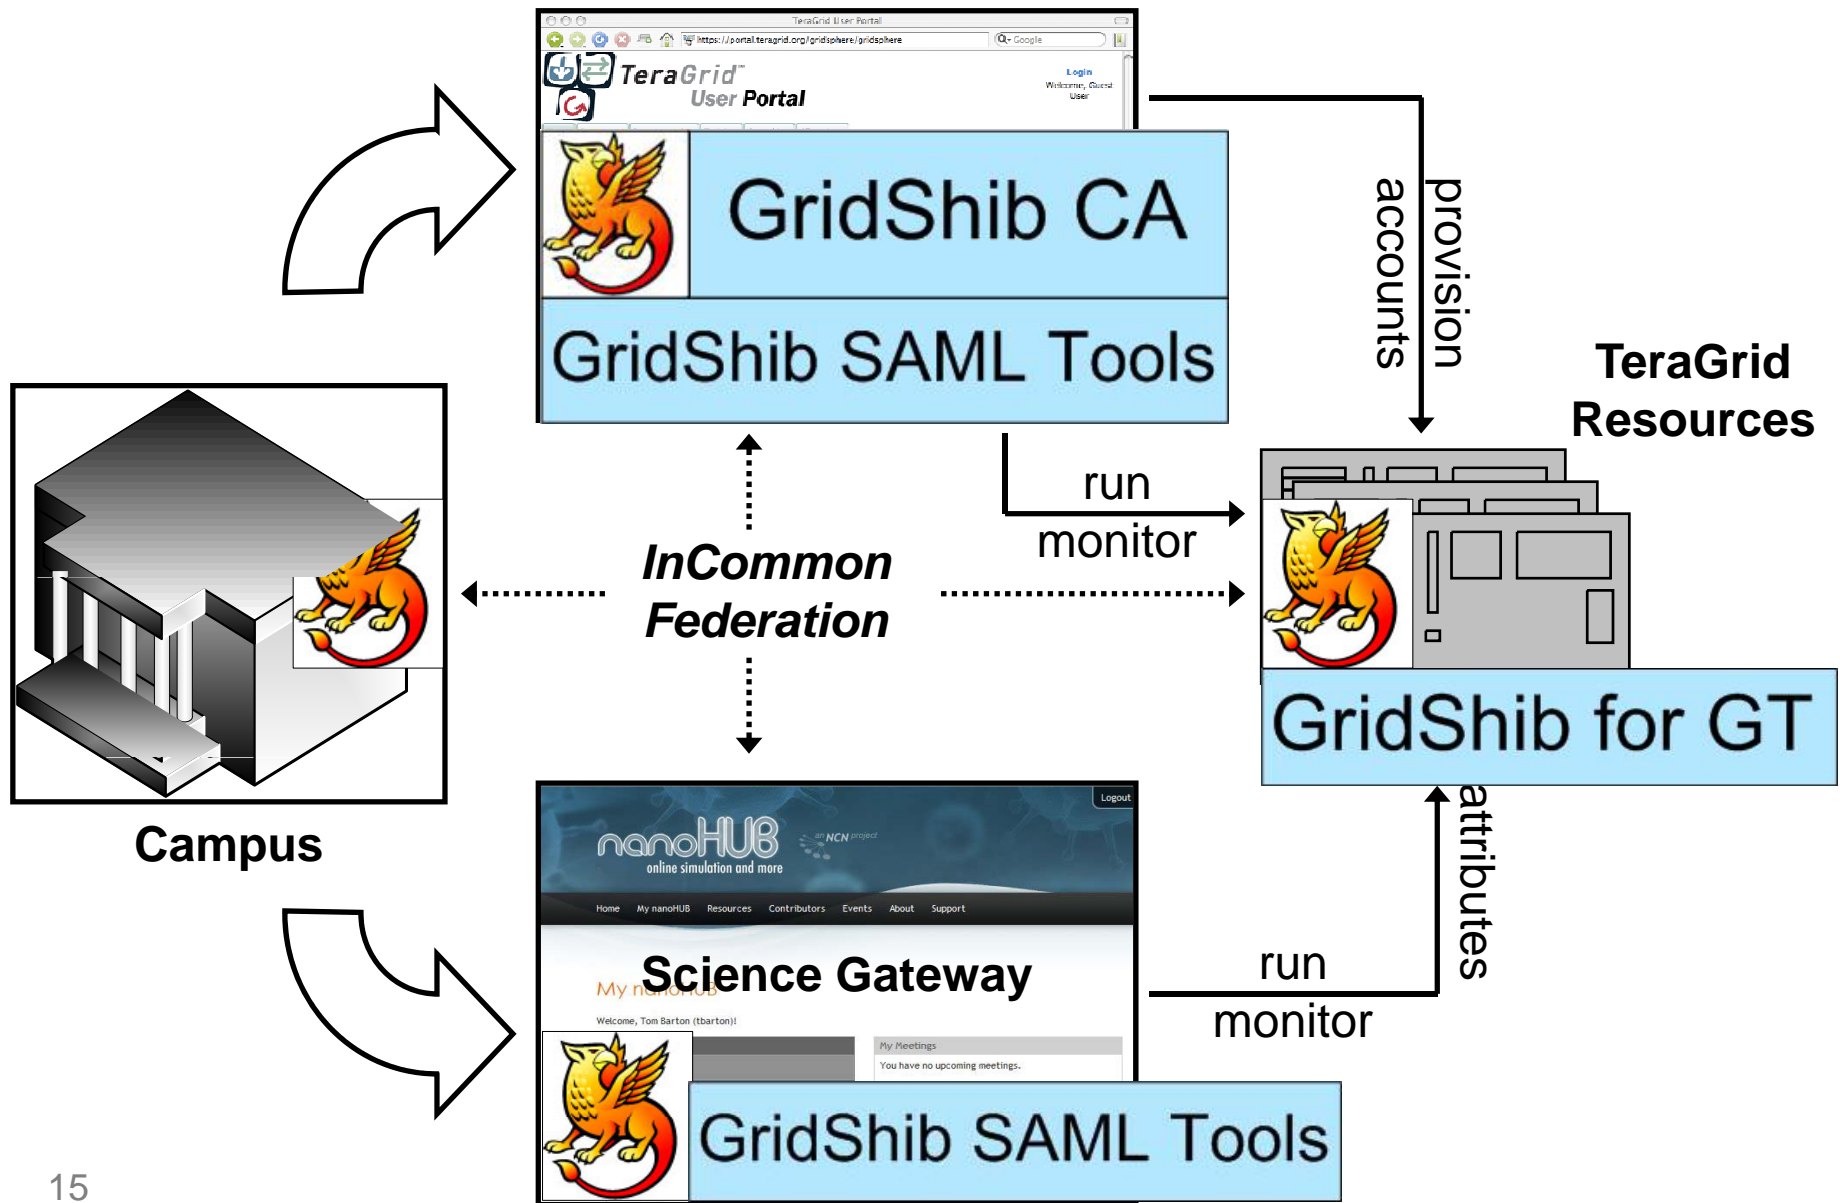
- Trust fabric: Metadata so that IdP's & SP's can mutually authenticate & interoperate
- Multilateral agreement among federation participants
 - Agree to actually operate as they claim to
- A “Where Are You From Service” available

TeraGrid Joining InCommon (as a Service Provider)

- Document high-level policy & procedure
 - What attributes are needed & why?
 - How are they handled?
- Agree to coordinate as necessary with other participants
- Status of privacy & security policies

GridShib Components





Managing Access

- Plenty of workable solutions for grids
 - VOMS, CAS, PERMIS, LCAS/LCMAPS, gJAF, SAZ, GridGrouper, ...
- Too many? Hinders grid-interop? VO-interop?
- Semantical & operational hurdles
 - Requires common semantics for attributes & groups, plus coordinated configuration of PDPs across resources, to yield consistent access practices



Managing Access

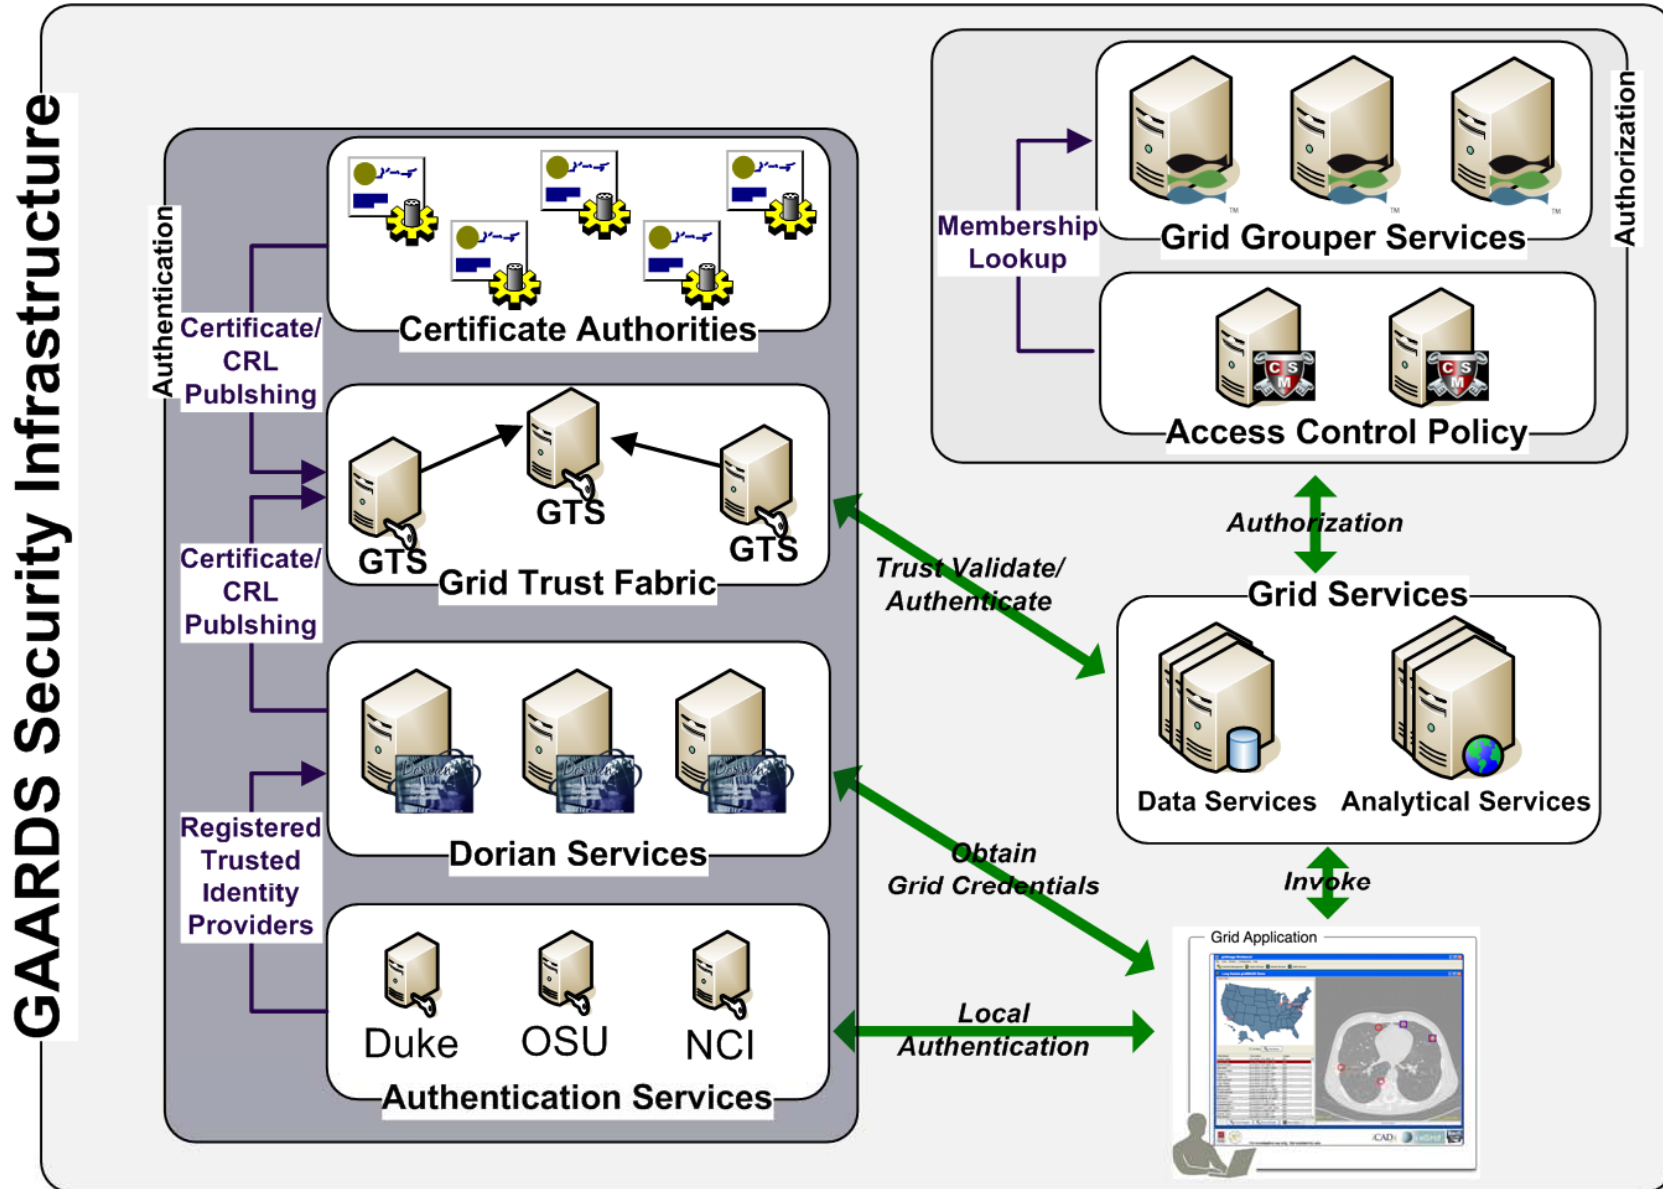
- Are all Sources of Authority integrated within a common access information management system?
 - No? One cause of proliferation of access management point solutions
 - No? Reduces transparency & auditability
- Are grid resources the only sort of value to collaborators?
 - Wiki, email lists, calendar, IM, video/audio conf, web presentation, webDAV, Course Management System, ...
 - These need their access managed, too

caBIG:

Cancer Biomedical Informatics Grid

- A “virtual informatics infrastructure that connects data, research tools, scientists, and organizations ...”
- **caGrid**: its underlying service oriented infrastructure
 - Local providers control access and management, but community accepted virtualizations of the data and analytics are made available using standardized service interfaces
 - caGrid v1.1 uses GT 4.0.3
- 50+ participating cancer research centers

caGrid's GAARDS



Grouper 101

- Groups are organized into Stems or Namespaces
 - URN-like names & delegation model
- Groups have members, including subgroups
 - Direct & indirect membership
 - Composite groups (union, intersection, complement of other groups)
- Metadata & privileges for Stems & Groups
 - Several delegation models for connecting Sources of Authority
 - Decorate groups with attributes
- Largest implementations to date have $O(10^5)$ groups and $O(10^6)$ memberships

GridGrouper

- It's a web service
- Several forms of delegation plus group math enables all Sources of Authority to participate
 - E.g. solution of multi-IRB access problem previously unsolvable
- Any number of GridGrouper instances can operate in the grid
 - Each service or resource identifies GridGrouper groups for access policy
 - Each research group is free to use central GridGrouper or run their own

GridGrouper UI

The screenshot displays the GridGrouper UI, a web-based interface for managing groups. The main window is titled "Group Management Browser" and shows a tree view of the group structure on the left and a detailed view of a selected group on the right.

Left Panel: Group Management Browser

- Grid Grouper Service(s)
 - https://cagrid02.bmi.ohio-state.edu:8443/wsrf/services/cagrid/GridGrouper
 - Grouper Administration
 - Grid Grouper Administrators
 - Ohio State University
 - Department of Biomedical Informatics
 - Staff (selected)
 - Students
 - people
 - Faculty
 - Authz Test

Collaboration Management Platforms

- Management of collaboration a real impediment to collaboration, particularly with the growing variety of tools
- Goal is to develop a “platform” for handling the identity & access management aspects of many different collaboration tools & resources
- This also presents possibilities for improving and unifying the overall user experience as well as UI for specific applications and components

CManage

- Being developed by the Internet2 community supported in part by an NSF OCI grant
 - Parallel activities in the UK and Australia
- Open source, open protocol
- Common access management using Shibboleth, Grouper, and Signet ...
 - Identity, Groups, Privileges, Federated Access
- ... across all integrated applications!

COmanaged Applications

- Now: wiki, blog, email list manager, audio conf, web meeting, calendar, ...
 - More collaboration tools on the way
 - Typical application integration issues with COmanage – no new hurdles
- Soon: grid integration with shib-grid integration technologies
 - E.g. use GridShib SAML Tools to integrate GT4 with COmanage



Federated login

InCommon

Select your Home Organization

In order to access a Resource on host 'comanage.internet2.edu' you must authenticate yourself.

Remember selection for this web browser session.

  This WAYF service developed by SWITCH. The [SWITCH](#) Foundation operates the Swiss Education & Research Network which guarantees high-speed connectivity to the Internet and to science networks globally for the benefit of higher education in Switzerland.

COmanage identity & access management console on top, application frame below



CO-Manage

A Collaborative Organization Identity Management Service
This is a demonstration service. Direct questions to gettes@internet2.edu.

Manage your CO identity, privileges and group memberships (use the following CO-dependent services).

tbarton@uchicago.edu

[My Identity](#)

[My Privileges](#)

[My Groups](#)

[My Services](#)

Welcome back!

Thank you for returning to the Collaborative Organization (CO) of "Internet2". Some of the following information maintained by the CO comes from your organization as a member of the [InCommon](#) Federation.

UPDATE

('*' denotes required information)

<i>Full Name</i> *	<input type="text" value="Tom Barton"/>
<i>First Name</i> *	<input type="text" value="Tom"/>
<i>Last Name</i> *	<input type="text" value="Barton"/>
<i>Email Address</i> *	<input type="text" value="tbarton@uchicago.edu"/>
<i>Phone</i>	<input type="text"/>
<i>Description</i>	<input type="text"/>

Group Member of i2nlr:mwsec:mace:mace-dir

2 [Delete yourself](#)

Current CManage services



CO-Manage

A Collaborative Organization Identity Management Service

This is a demonstration service. Direct questions to gettes@internet2.edu.

Manage your CO identity, privileges and group memberships (use the following CO-dependent services).

tbarton@uchicago.edu

[My Identity](#)

[My Privileges](#)

[My Groups](#)

[My Services](#)

The following services *may be* available to you

[Mailing Lists](#)

[Confluence](#)

[MACE-Dir Group Protected Page](#)

[Join a Web Meeting \(dimdim\)](#)

[Blogging \(WordPressMU\)](#)

[Manage Voice Conferences \(asterisk\)](#)

[Bedework Calendar](#)

Federated Calendar	Administration	Public Event Calendar
------------------------------------	--------------------------------	---------------------------------------

You may opt-in or opt-out of the group

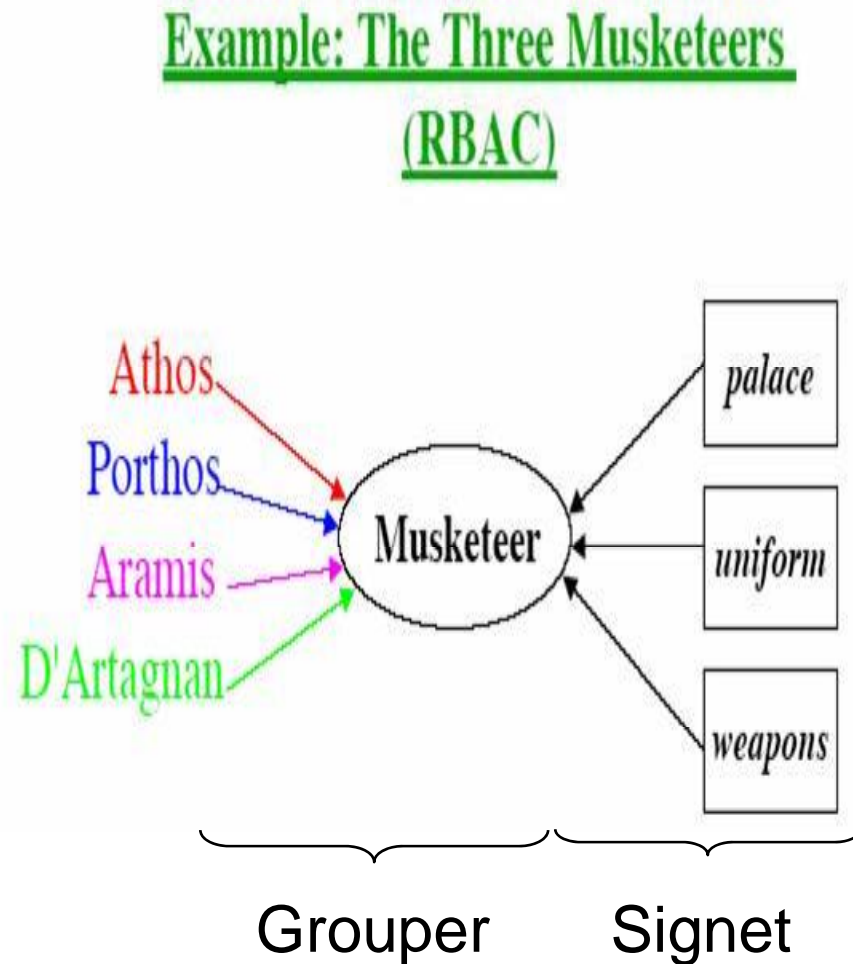
Internet2 Collaborative Organization: Middleware and Security: MACE: [MACE-Dir]
in *My Groups* above.

The task responsible for reflecting group data into the CO directory runs every 60 seconds. There may be an additional delay until CO services notice the changes in the CO directory.

Relative Roles of Signet & Grouper

RBAC model

- Users are placed into groups (aka “roles”)
- Privileges are assigned to groups
- Groups can be arranged into hierarchies to effectively bestow privileges
- Grouper manages, well, groups
- Signet manages privileges
- Separates responsibilities for differing Sources of Authority



Privilege Elements by Example


By authority of <i>the Dean</i>	grantor
<i>principal investigators</i>	grantee (group/role)
who have completed <i>training</i>	prerequisite
can <i>approve purchases</i>	function
in the <i>School of Medicine</i>	scope
for <i>research projects</i>	resource
up to <i>\$100,000</i>	limit
until <i>January 1, 2009</i> as long as <i>a faculty member at...</i>	conditions



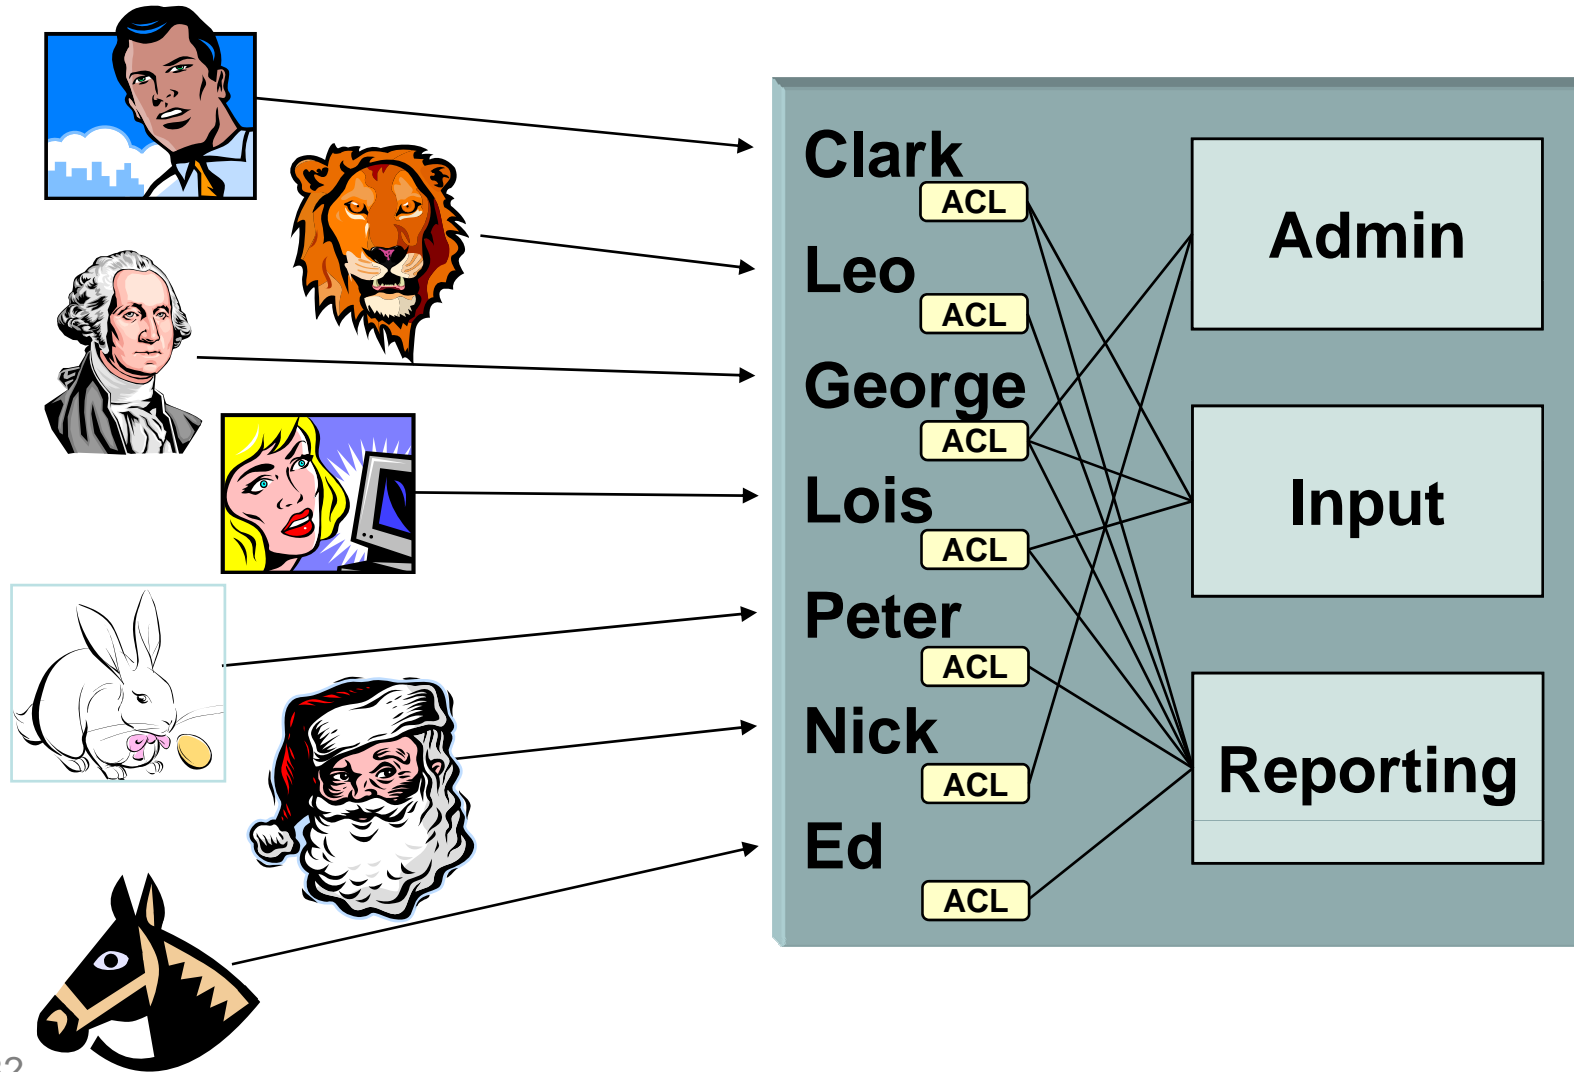
Semantics & policy, again

- Common semantics for attributes & groups plus common practice in configuring PDPs yield desired access practices
 - Hard, slow, unenforceable. Problems only detectable by use
 - We're comfortable with groups, which leaves us with semantical problems that must be solved outside of our management tools
 - And we don't know what that "priv" stuff is all about
- Distributed authority management might at least provide a framework to address some of the semantical problems...

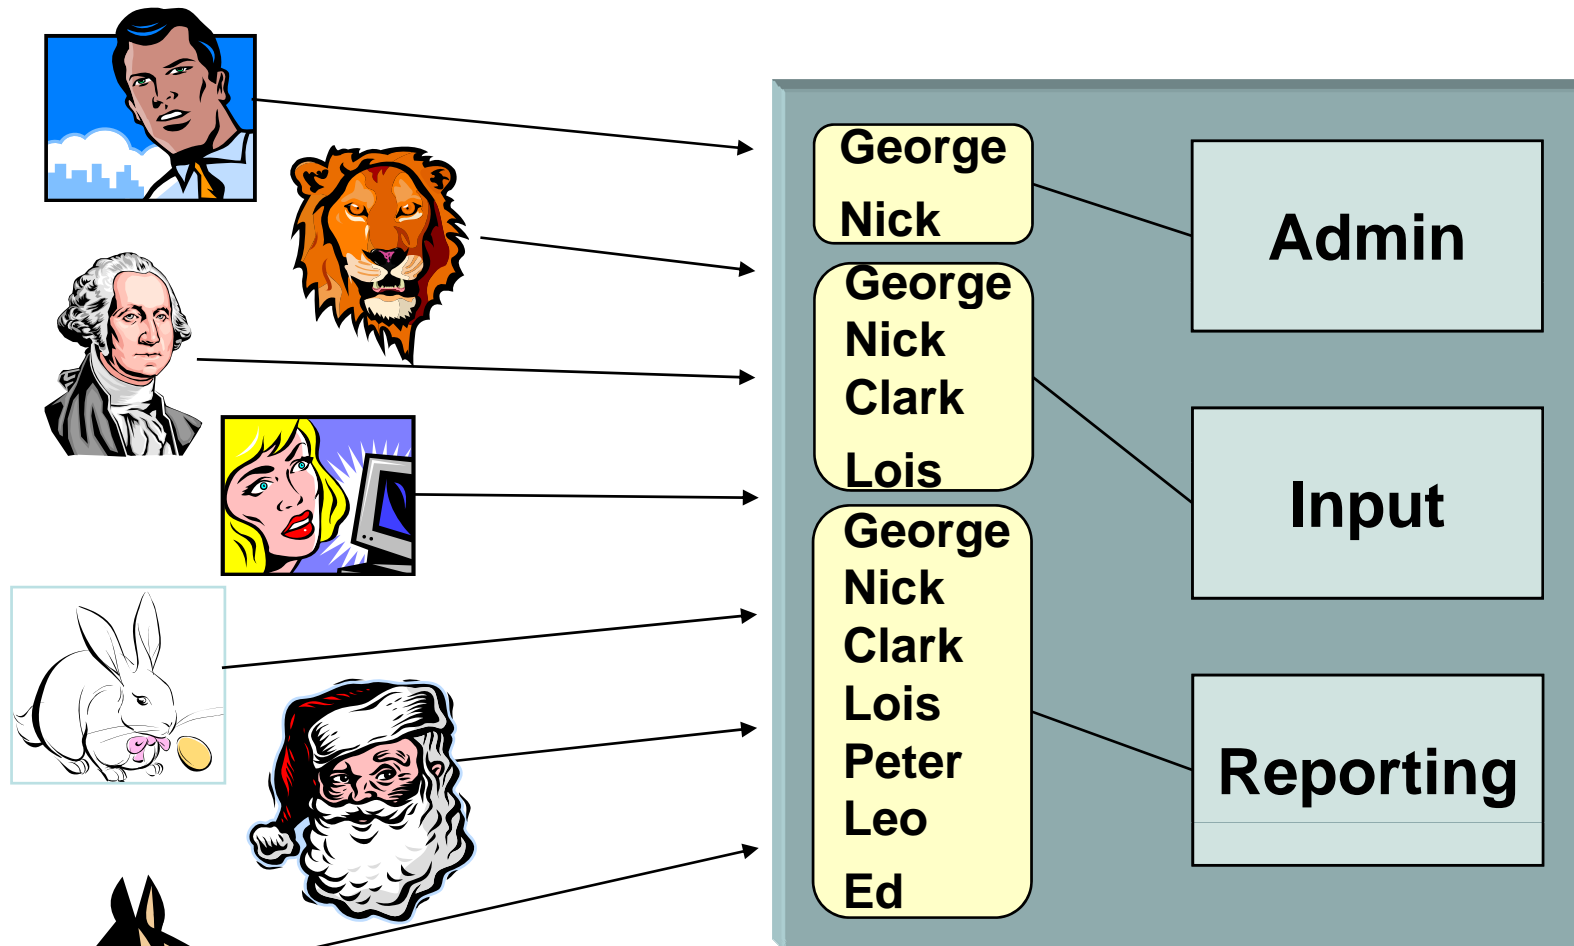
Configured how?



Stone Age

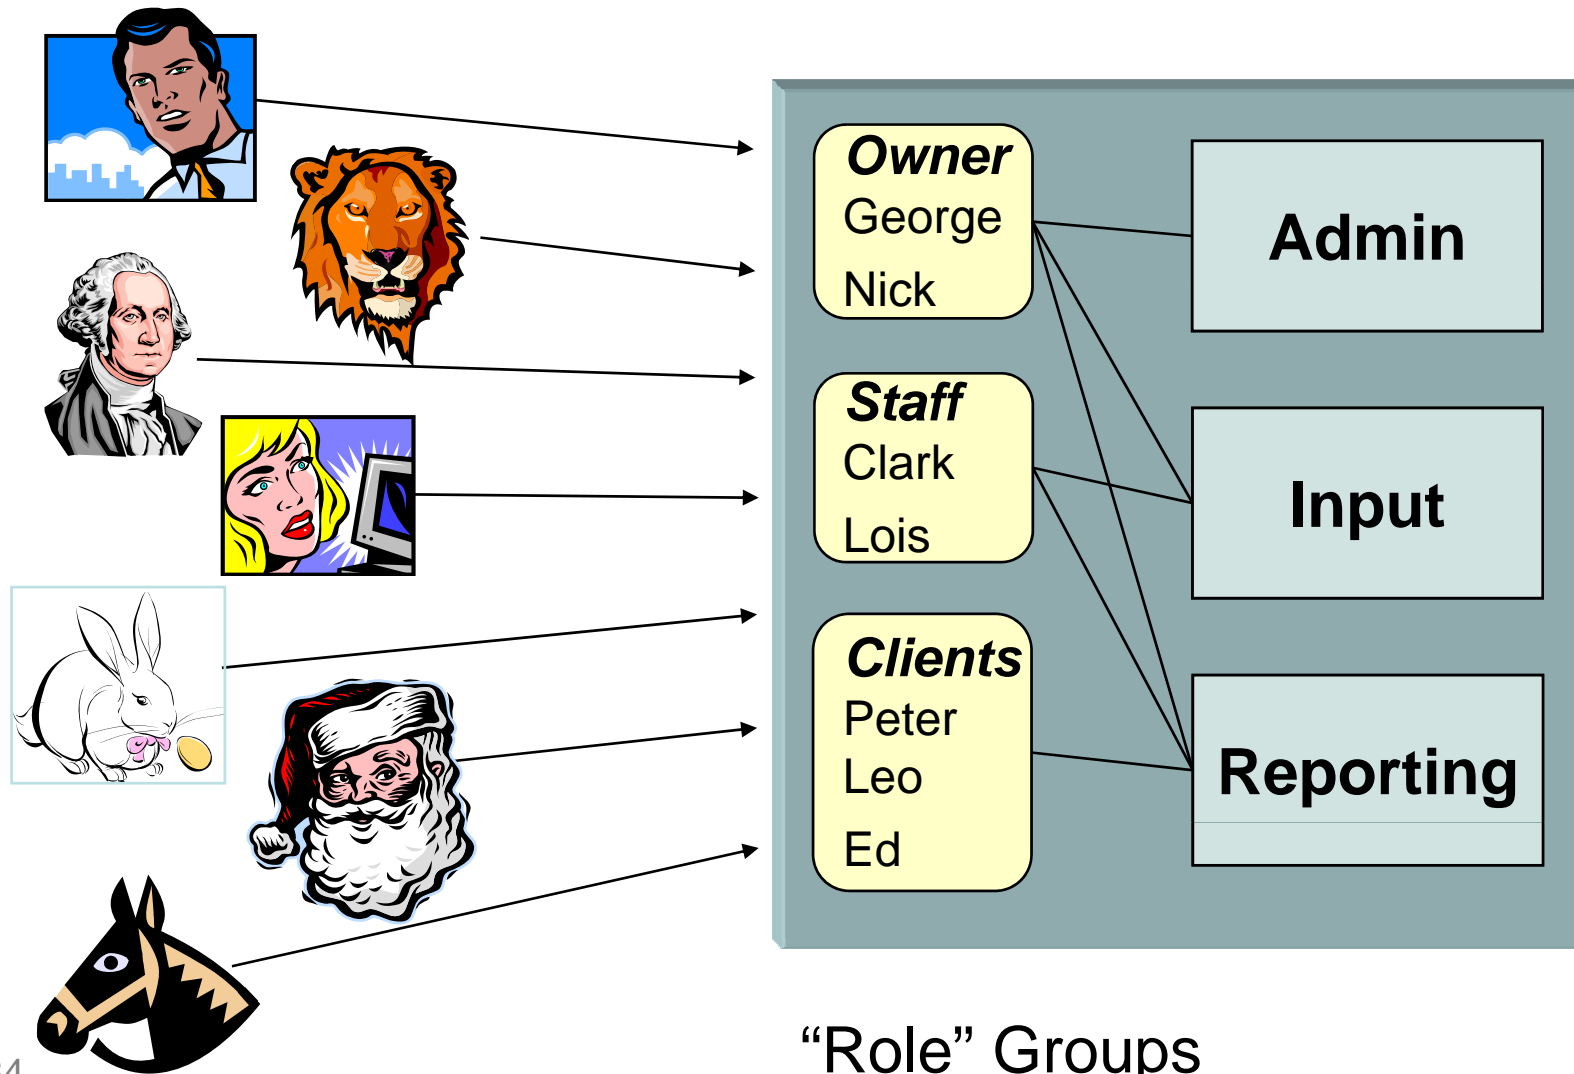


Middle Ages



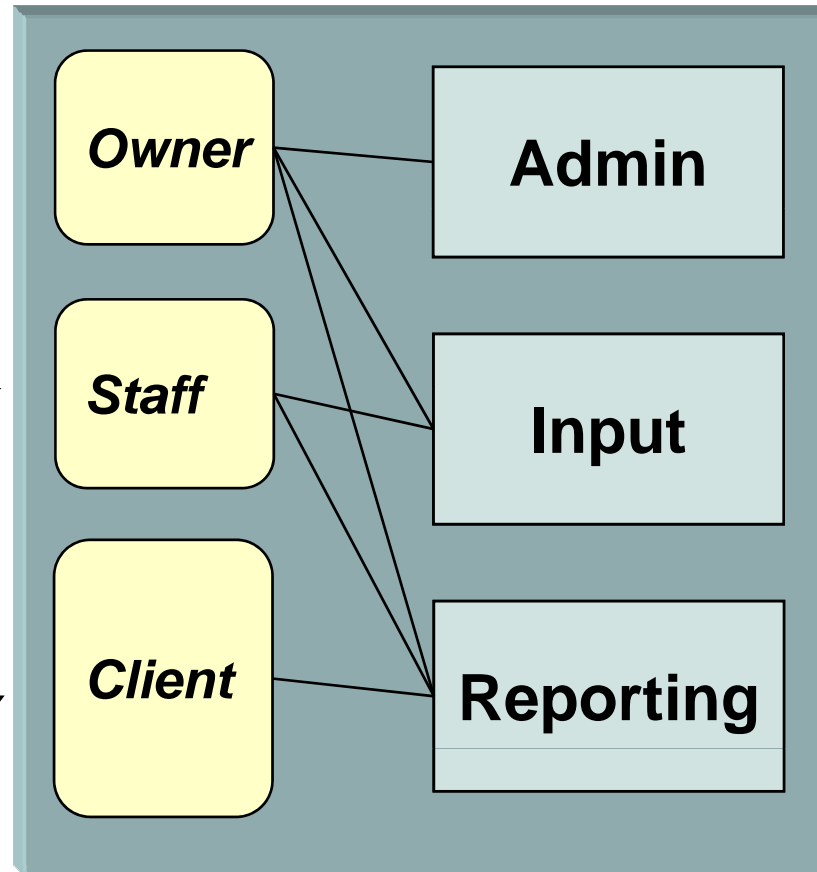
Functional Groups

Renaissance



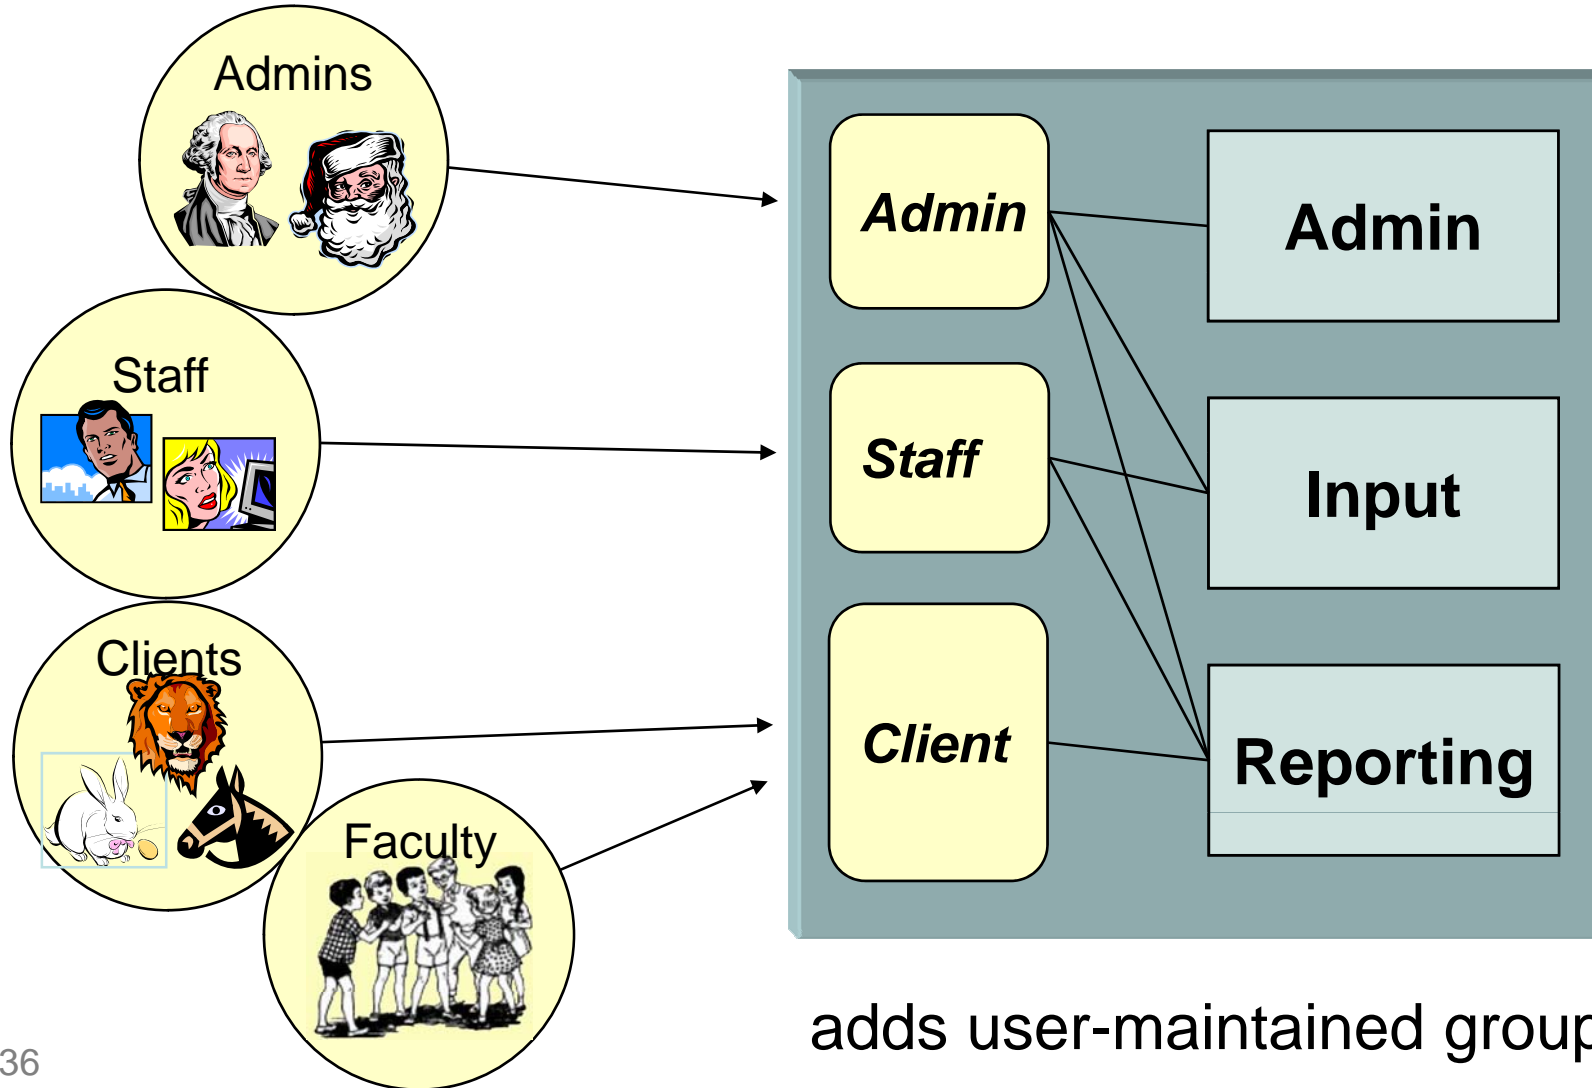
20th century

Identity Management!




Enterprise roles, affiliations

Groups Management

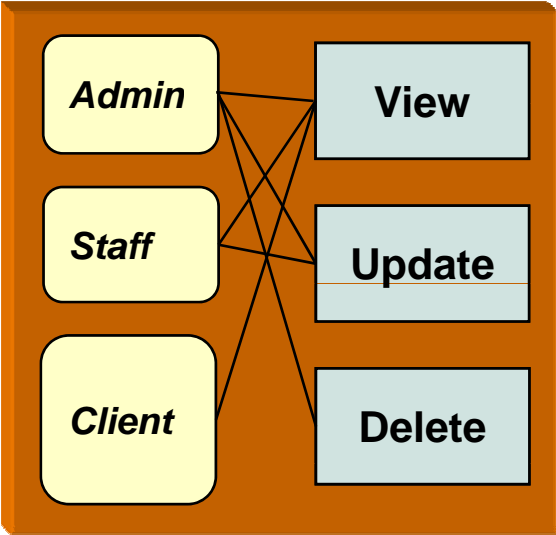


adds user-maintained groups

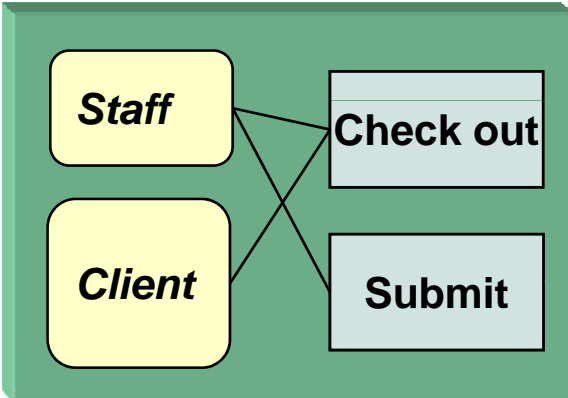
Something still missing



Each system ...



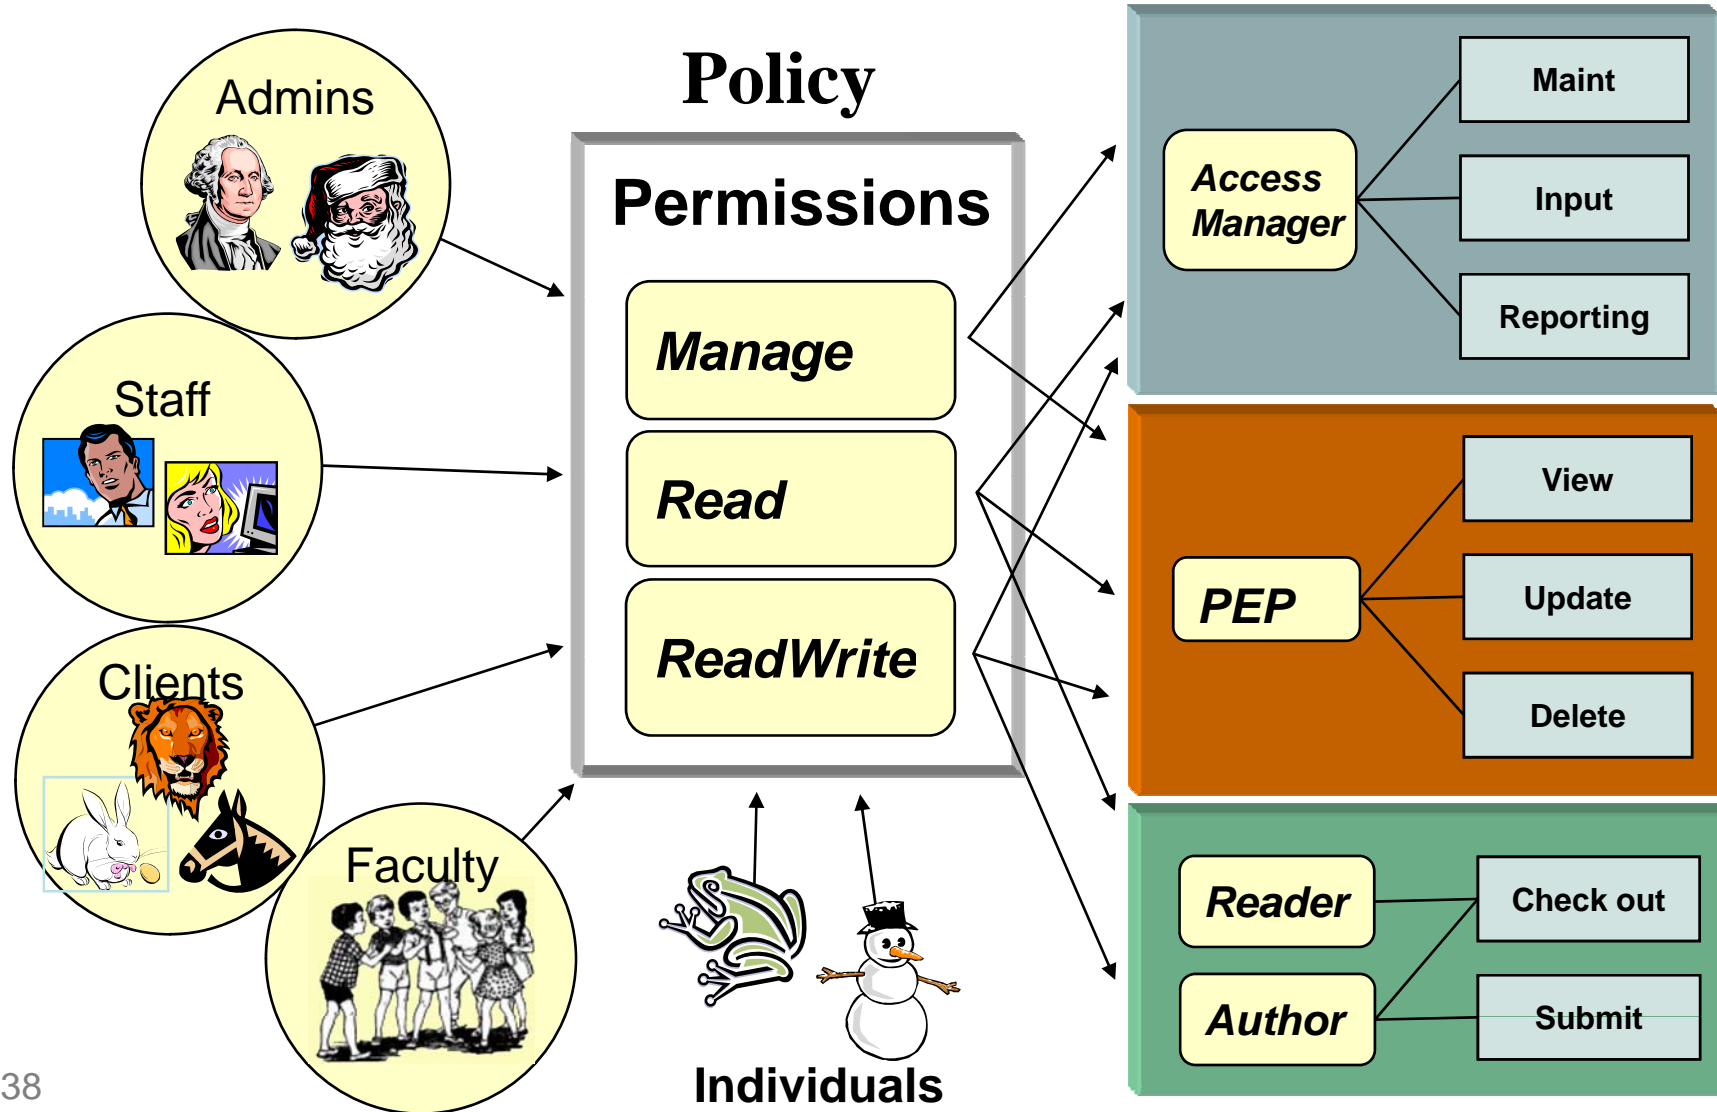
interprets policy ...



and sets access rules ...

separately.

Privilege Management



An Example: Stanford's Authority Manager

- Divisions, units, departments do not operate alike
 - A university-wide access policy based on roles cannot succeed
 - How about EGEE, OSG, (EGEE U OSG) –wide?
- Their solution: Distribute the authority for managing access to a unit's stuff to those responsible for the unit & integrate application security with Authority Manager
- $O(10^4)$ different privileges assignable
- $O(10^5)$ privileges assigned
- Internet2's Signet is derivative of Stanford's Authority Manager