Internet-Scale Identity and Collaboration

Dr. Ken Klingenstein,
Senior Director,
Internet2 Middleware and Security

Topics

- The Rise of the Internet
  and the Rise of the Middle layer
- Internet-Scale Identity
- Collaboration Tools
- Putting the Parts Together
  - Especially for VO’s
  - Some principles for application designers

The Rise of the Internet

- Making the technology: The late sixties and seventies established the core TCP/IP technologies and value to the CS community
- Making the market: The eighties made a mass market of technology, applications and content
- Making the business: The nineties created business plans and businesses

Takeaways

- Modular and layered design
- Open standards, open source
- Autonomous systems, loosely coupled
- Making a market is critical; sales are hard at first

The Rise of the Middle layer

- Development of campus and enterprise services common to many applications
  - Directories, enterprise authentication, group and privilege management, identity management
  - Policies and business processes
- Federations to extend middleware to interrealm needs
  - Attributes and federating software
  - Trust policies

Where we are now

- Many enterprises have basic middleware services connected to some applications
- Federations, and federated identity, are growing and learning to interact with other federations
- We are close to resolving Internet identity
- We are just beginning to understand linked identities, attribute flows, privacy mechanisms, etc.

Takeaways

- Modular and layered design
- Open standards, open source
- Autonomous systems, loosely coupled
- Making a market is critical; sales are hard at first

Requirements for Internet identity

- Fewer Internet sign-ons
- Preservation of privacy, especially across international boundaries
- Several layers of assurance of identity, to deal with low-risk to high-risk applications
- Ease of deployment
- Ease of use

Types of Internet identity

- Federated
  - Leveraging enterprise identity for inter-realm purposes
  - Authentication, entitlements and attributes are the common payloads
  - Privacy, security and trust are the critical issues
  - Is hard to do
- P2P
  - Originally PGP, now Infocard, OpenId, etc.
  - Need trust fabrics - may be coupled with reputation systems for trust
  - Is easy to do
- Both are growing at exponential rates

Federated Identity

- Enterprises exchanging assertions about users
  - Often identity based but can provide scale and preserve privacy through the use of attributes
  - Real time exchanges of standardized attribute/value pairs
- Basis for trusting the exchanged assertions via common policies, legal agreements, contracts, laws, etc.
- Federations offer a flexible and largely scalable privacy preserving identity management infrastructure

An adoption curve

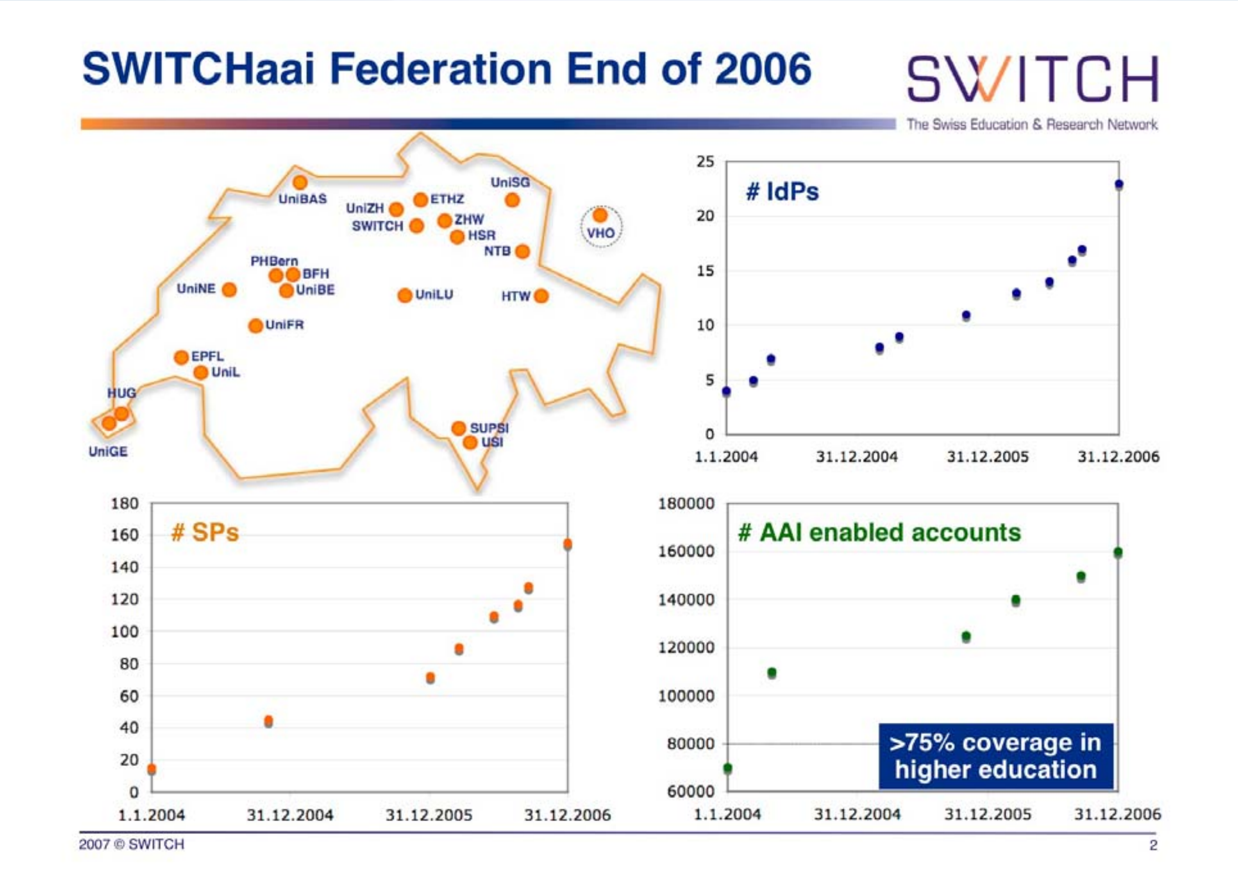

SWITCHaai Federation End of 2006

SWITCH
The Swiss Education & Research Network

>75% coverage in higher education

2007 © SWITCH

2

Presenter’s Name

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 bi-lateral; almost all in the R&E world are multilateral
- They provide a powerful leverage of enterprise credentials
- Federations are learning to peer
- Internal federations are also proving quite useful
- (Note: federated *.*, but not Internet identity scale)

International R&E federations

- Substantial deployments in many countries, including UK, Norway, Switzerland, US, Australia, France, Denmark, Finland, Spain, Germany, Netherlands, etc.
- Most are Shib based; some use other SAML products, PAPI, etc…
- Scope of membership usually higher ed, but some are broader, e.g. UK, Spain, Netherlands
- Use cases range from content access to collaboration support to learning management systems to wireless roaming to…
- Many are NREN-leveraged; some like IGTF are not

Technical Aspects of Federations

- Federating protocol
- Enterprise signing keys
- Metadata management and WAYF service
- Enterprise Identity Management practices

Policy Aspects of Federations

- Participant operational practices
- Agreement between federation and members
- Standardized attributes
  - eduPerson
  - Levels of Assurance (LOA)

Relationships among federations

- Peering
- Confederation
  - Presumes peering, adds multi-federation support
- Leveraged
  - Specialized federations that extend a common base federation – e.g. the California system
- Intersecting

Peering Parameters

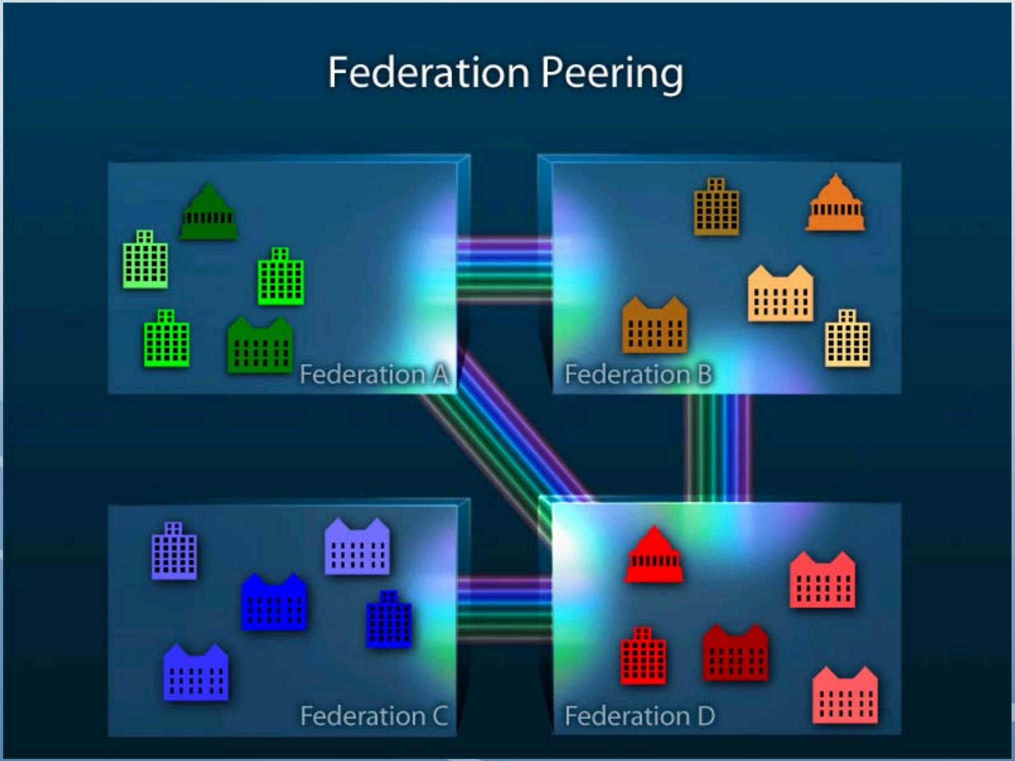

Parameters:

- **LOA**
- **Attribute mapping**
- **Legal structures**
- **Liability**
- **Adjudication**
- **Metadata**
  - **VO Support**
- **Economics**
- **Privacy**

Some inter-federation key issues

- Multi-protocols
- Sharing metadata
- Aligning policies
- WAYF functionality
- Dispute resolution
- Virtual organization support

Prague Meeting on Inter-federation

- 15-20 International R&E federations
- Hangers on: Liberty Alliance, ITU, Gartner, etc…
- Several key areas for agreements - LOA mapping (generally okay), Attribute mapping, Privacy Policies, Dispute resolution, Financial considerations, Technical direction setting
- Ongoing process mechanism
- Prague, September 3

P2P Identities

- Provides tokens for interpersonal trust
- Initially PGP, now OpenId, Infocard
- Use cases include blogs and wikis, file and photo sharing, some encrypted email, etc.
- Active space – Cardspace in MS Vista, Higgins and the Bandits, OpenId, etc.
- Several layers
  - Globally unique identifier
  - Hooks to a trust or reputation system
  - Mobility solution
  - Protocol layers

P2P Development

- Growth is dramatic
- Plugs into almost any application
- Integration with Infocard
- Starting to hit the hard issues:
  - Revocation
  - Delegation and transitive trust
  - Privacy

Identity integration goals

- First, of federated and p2p identity
  - Many levels of integration – tokens, GUI, privacy management paradigm, trust fabrics…
- Then, of identity and privilege management
  - Assignment and management of permissions to users by those with authority to grant such access
  - Addresses the static aspects of the authorization space, with audit, delegation, prerequisites, etc.
  - Permissions can be enterprise or virtual organization

A Bloom of Collaboration Tools

- An over-abundance of new tools that provide rich and growing collaboration capabilities (aka Web 2.0)
- Do you
  - Wiki, blog, moodle, sakai, IM, Chat, videoconference, audioconference, calendar, flikr, netmeeting, access grid, dimdim, listserv, webdav, etc
  - Share files among workgroups, access Elsevier, work with the IEEE, etc
- No uber-app – limits invention and community of users
- 3 - 4 is fine, but many per user is hard to manage

Collaboration Tools design issues

- Asynchronous vs synchronous
- Integration of content across tools
- Managing presence
- Community of use
- Managing privacy
- Many, many tools have overlapping parts that do not interoperate
- How much collaboration can we handle?
  - People
  - Tools

Collaboration Tools and Identity Management

- Deeply enriches collaboration tools
  - Fine-grain access control and wikis
    - spaces.internet2.edu
    - “member of the community” processes
  - Transparently shared file stores
  - Collaboratively visible calendaring
  - Embedded VO IM channels in campus portals

Relieving the Pain of Rich Collaboration Management

- Commonly manage which identities and which attributes can use the capabilities of the collaboration tools
- Can offer delegation, privacy management, maybe even diagnostics
- COmanage

Collaboration Tools and Identities

- Enterprise, VO, and P2P persona are in all of us – our day job, our second job, the rest of our life…
- When and how we integrate the persona needs to be carefully done – legal, ethical, personal issues
- The abundance of communication and collaboration devices makes this harder

Putting It All Together

- Real life and the attribute ecosystem
- Identity management and VO’s
    - The collaborative processes
        - “Internet-scale collaboration”
    - The domain work
- Takeaways for developers

Real life and the attribute ecosystem

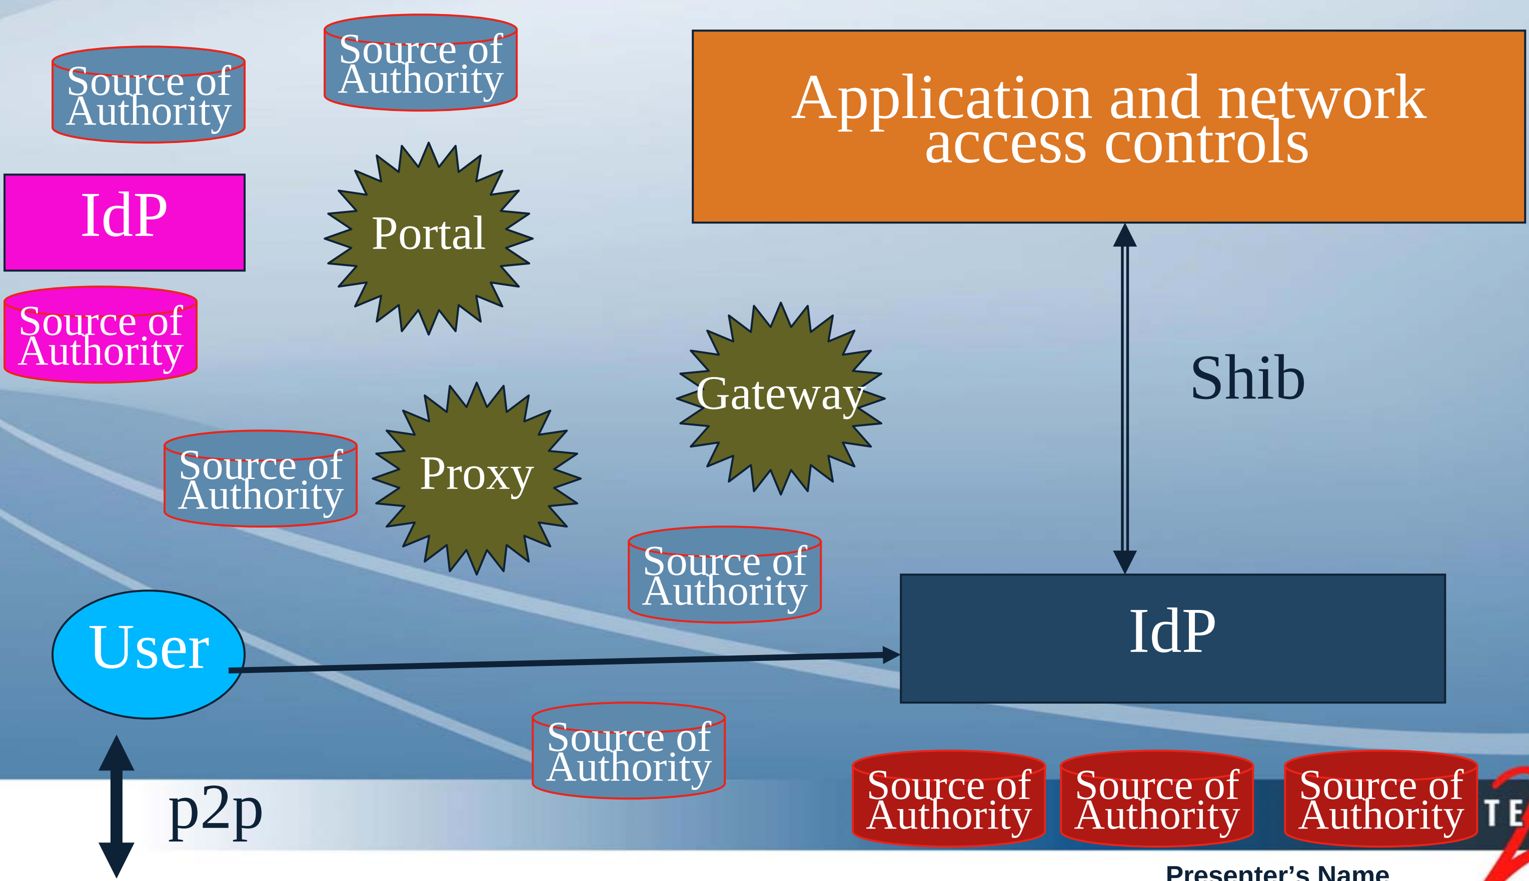

Source of Authority

Source of Authority

Application and network access controls

IdP

Portal

Source of Authority

Shib

Gateway

Source of Authority

Proxy

Source of Authority

User

IdP

Source of Authority

p2p

Source of Authority

Source of Authority

Source of Authority

Identity Management and VO’s

- Enterprises, federations, and peering can offer VO’s flexible and sustainable identity management support
  - Collaboration
  - Domain science
  - Integrating instruction and research
- Commitments seem to be there, and the technologies are being deployed…
- Transitions from current duct-tape can be done incrementally

VOs plumbed to federations

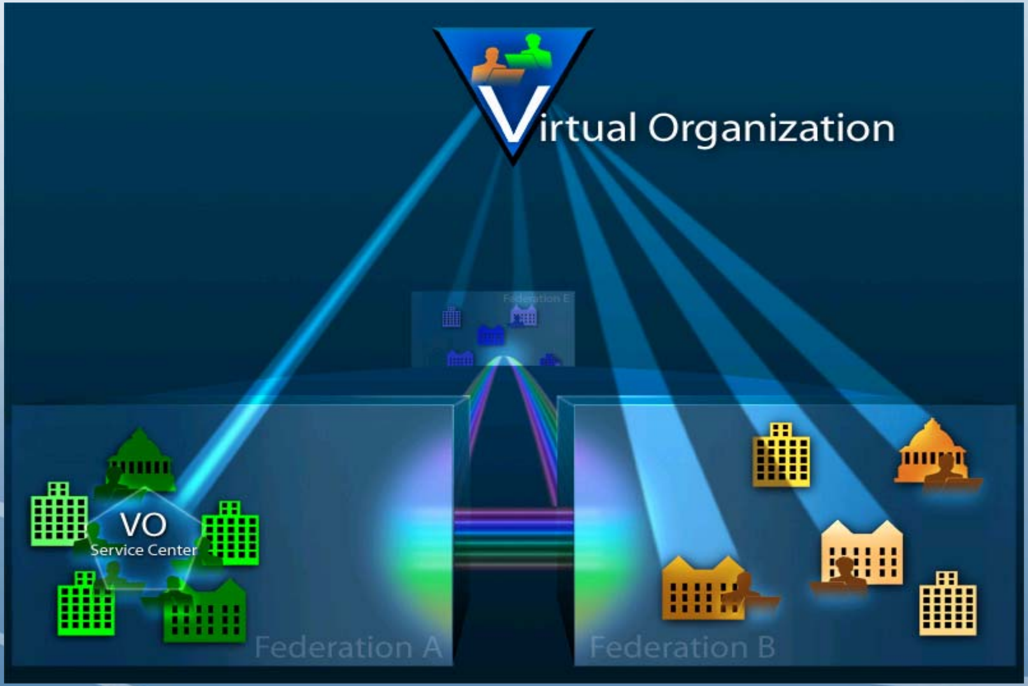

Collaboration and Virtual Organizations

- VOs are first collaborative organizations
  - General collaboration tools – listservs, wikis, audioconferencing, videoconferencing, shared calendars, etc.
  - Academic collaboration tools – grant proposal and administration management, paper development and publication
- Many support components for such activities can also meet needs in the domain science management

Domain science and federations

- Identity, management of groups, management of privileges to those identities and groups in the domain applications can be shared with the collaboration piece
- First steps are federated identity:
  - Several projects are combining Grids and Shibboleth
  - A bio-informatics grid is using federated identity and some group management
  - Most are entry level work designed for quick benefits
- Final steps lie in attribute-oriented controls and functioning within the attribute ecosystem
- Other aspects, such as workflow and diagnostics may be integrated

Integration of education and research

- Class lists can be easily assigned domain science capabilities; TA's can have simple domain privileges
- Domain science materials can be presented within collaborative settings
- Scientific, scholarly and business workflows can interoperate

Getting there from here

- Open standards underlie the work
- VO Service Centers can play multiple roles, from collaborative service platform instances to training others on deployments of platforms
- Ultimately the challenge is about the applications depending on infrastructure more than the management platform itself
- Existing applications, especially big ones, are hard to reengineer

Takeaways for application developers

- Understand the attribute ecosystem
- Leverage the trust fabrics
- Properly done, the infrastructure can bring in lots of use cases to the domain.
- Be conservative in the data you send, be liberal in the data you accept
- The first thing one learns from an interoperability protocol is all the ways in which we can’t operationally interoperate

A few more takeaways

- The sooner you start, the longer it takes
- Try doing it with the engine running
- Perfection is achieved, not when there is nothing more to add, but when there is nothing left to take away.
- The only numbers of importance in computing are 1, 2 and many - with its meta counting variant: 1, 2, Schema
- Any piece of software reflects the organizational structure that produced it
- In theory, there is no difference between theory and practice; In practice, there is