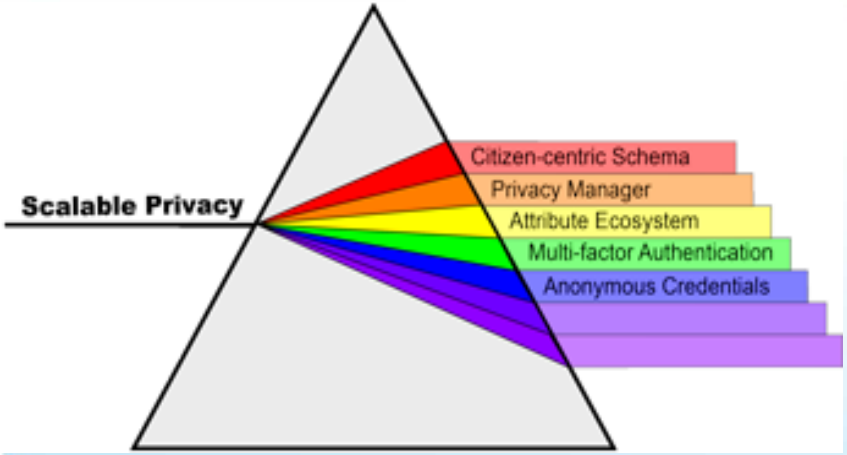

September 2013
mgrady@unicon.net
Michael Grady, Unicon, for the Internet2 Scalable Privacy Project

Scalable Privacy and Multi-factor Authentication Pilot

Context: InCommon today

- 400+ universities, 650+ total participants, growth continues rapid
  - Many cloud service providers, from Microsoft to Elsevier to NIH and NSF to ***
- > 10 M users
- Traditional uses continue to grow:
    - Outsourced services, government applications, access to software, access to licensed content, etc.
- New uses bloom:
    - Access to wikis, shared services, cloud services, calendaring, command line apps, medical, etc.
    - A basic requirement for cloud services
- FICAM certified at LOA 1 and 2 (Bronze and Silver).
- New services
    - Certificates – SSL and Personal
    - InCert - open-source client-cert lifecycle management
    - Certification marks - R&S (Research and Scholarship)

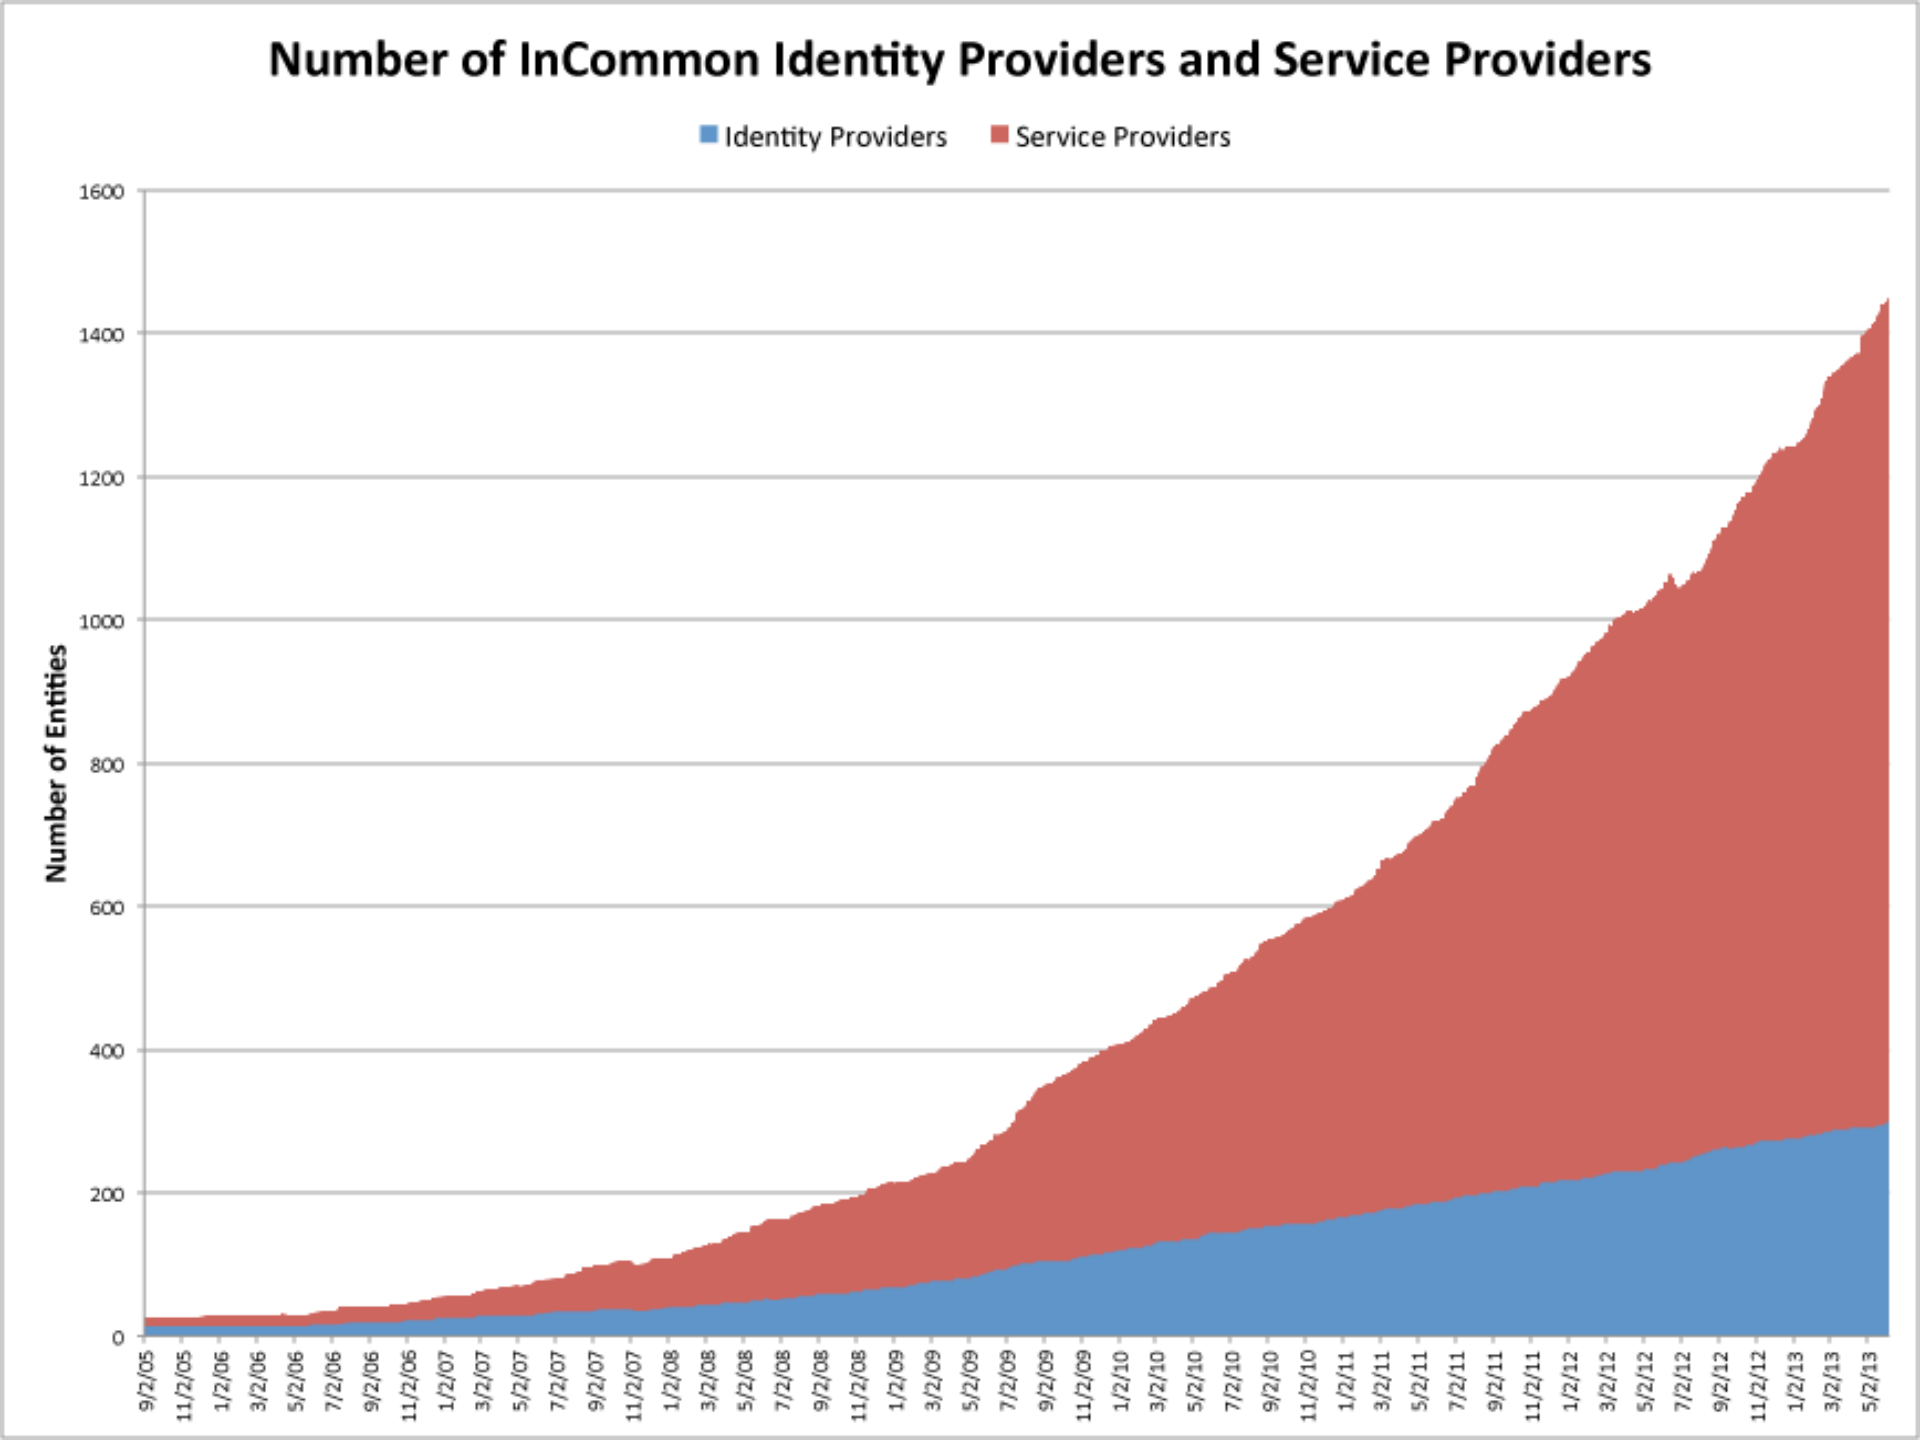
Number of InCommon Identity Providers and Service Providers
Identity Providers
Service Providers
Number of Entities
0
200
400
600
800
1000
1200
1400
1600
9/2/05
11/2/05
1/2/06
3/2/06
5/2/06
7/2/06
9/2/06
11/2/06
1/2/07
3/2/07
5/2/07
7/2/07
9/2/07
11/2/07
1/2/08
3/2/08
5/2/08
7/2/08
9/2/08
11/2/08
1/2/09
3/2/09
5/2/09
7/2/09
9/2/09
11/2/09
1/2/10
3/2/10
5/2/10
7/2/10
9/2/10
11/2/10
1/2/11
3/2/11
5/2/11
7/2/11
9/2/11
11/2/11
1/2/12
3/2/12
5/2/12
7/2/12
9/2/12
11/2/12
1/2/13
3/2/13
5/2/13

Scalable Privacy

- 2+ year grant to Internet2/InCommon
- Development partners are CMU, Brown, with expertise from Wisconsin, Ohio State and others
- Several focal points
  - Promotion of multi-factor authentication
  - Citizen-centric attributes and schema
  - Development and deployment of privacy managers
  - Introduction of anonymous credentials
- https://spaces.internet2.edu/display/scalepriv

Work described in this presentation is supported by the National Strategy for Trusted Identities in Cyberspace (NSTIC) National Program Office and the National Institute of Standards and Technology (NIST). The views in this presentation do not necessarily reflect the official policies of the NIST or NSTIC, nor does mention by trade names, commercial practices, or organizations imply endorsement by the U.S. Government.

INTERNET 2

Key deliverables

- Promotion of two factor authentication
  - Good privacy begins with good security
- Citizen-centric attribute activities
  - For transactions, for accessibility, for social government
- Trusted metadata approaches
  - About the relying party and the Identity Provider
  - Vetted by the federation and by third-parties
- Next-generation privacy manager
  - Leveraging prior work, trusted metadata, usability-built-in
- Anonymous credentials
  - Evaluate issues in integrated deployments at scale
  - Integration in software, use of metadata, and user experience

INTERNET 2

How it all fits together

- A user, in their context as a university student, uses a privacy manager to release their institutional affiliation to student discount services
- A user, in their context as a citizen, uses a privacy manager to release sufficient residence information that allows them to then anonymously post to the neighborhood-only wiki.
- A user, in their context as a consumer, uses a privacy manager to manage the release of preferences (e.g. zip code, preferred language, geolocation, etc) to customize commercial services while preserving privacy
- With your paper diploma and your identity-rich e-transcript, you get issued an anonymous token asserting affirmation of graduation and degree, year, honors, major
- A user, in their context as a worker, uses a privacy manager to release anonymous credentials (such as security clearances and personal medical information) to third party contractors.
- A parent uses a privacy manager to manage their children's on-line privileges to COPPA-compliant applications

INTERNET 2

Promotion of multi-factor authentication (MFA)

- Good privacy begins with good security
- MFA addresses a significant number of security threats
- A variety of second factor alternatives are now viable – USB devices, NFC devices, cell phones, certificates, etc., and technology can bridge across them
- Advantages of MFA and Federated identity
  - Combining MFA with WebSSO and federated identity allows MFA to be leveraged by many services/SPs; "MFA externalities"
  - Potential to help achieve higher levels of assurance
  - If biometric factors are used, "privacy spillage" limited to IdP

INTERNET 2

MFA: Two major thrusts

- MFA Pilot Institutions: support wide-scale deployments of MFA technologies at three institutions:
  - Massachusetts Institute of Technology (MIT)
  - University of Texas System
  - University of Utah
- MFA Cohortium: Create and facilitate a cohort of additional institutions, establishing a collaborative environment for sharing questions, requirements, planning, expertise, experience, artifacts, etc. related to deploying and supporting MFA, leveraging the pilot institution activities.
  - Now ~ 40 institutions, > 1M potential users
  - Creating a next generation of MFA aware users
  - Technology agnostic, lifecycle oriented
  - https://spaces.internet2.edu/display/mfacohortium

MFA Cohortium

- Mix of institutions wondering about MFA, starting to deploy MFA, and a few with reasonably significant deployments
- Help campuses without MFA understand the need for it, the risks it addresses, its costs, etc.
- Help campuses that are implementing MFA with deployment, policy, technology, usability and accessibility
- Collect and create extensive set of resources/artifacts on “all things MFA planning and deployment” for Higher Ed, establishing a public web site to serve as lasting resource site. https://spaces.internet2.edu/x/4AwwAg
- Four subgroups formed to begin executing work plan: Business Case, Deployment Strategies, Technology, Product & Vendor topics

Early interesting issues in MFA at scale

- Accessibility support
  - From device issues to accessing preferences during MFA processes
- FERPA issues in the release of PII (e.g. cell phone number) to third party authenticator
  - More generally the legal relationship between enterprise and third party authenticators
- Cloud authenticators and DDOS attacks
  - Should enterprise authn fail under external DDOS?
  - Generally, identify key barriers to outsourcing components of authn
- Alternative strategies when multifFactor tokens aren't available
  - MFA fails more frequently, if only for environmental issues
  - "Fallback" approaches for opt-in deployment models?
- ROI of federated MFA
  - The leverage of federation and MFA is enormous, but how to capture it?

INTERNET 2

Three important software deliverables

- Shibboleth-based integrated, universal MFA handler
  - Shib is the most widely used open source federating software platform in the world
  - Multilateral Shib-based federations exist in over 40 countries, in real estate, in government, in law enforcement, in securities and banking, etc
  - A universal well-integrated MFA handler instantly opens MFA externalities
- CAS integrated, universal MFA handler
  - CAS is a very widely used open source SSO
- InCert
  - Open source client certificate lifecycle management system
  - Also provides device boarding and device security
  - Client certs are invaluable for many ecosystem capabilities beyond authentication and anti-phishing
  - http://www.internet2.edu/incert/
  - https://spaces.internet2.edu/x/vAhOAg

Key MFA artifacts (some ways along)

- Living list of Issues Identified & Lessons to Learn
- Alternative Strategies When Multi-Factor Tokens Are Not Available
- Initial Deployment Strategies for Multi-Factor Authentication
- Deployment Decision Tree diagrams
- Multi-Factor Authentication Solution Evaluation Criteria

Key MFA artifacts (starting)

- Architectural patterns of integration
- Technology-focused documents:
  - Comparative analysis/review of the various security properties of MFA technologies: highlight strengths & weaknesses
  - Technologies Assessment Matrix: Assess on multiple factors, with broad categories such as Security, Usability, and Deployability, plus considering accessibility
- Business Case for Multi-Factor Authentication: identify elements, create template, and gather examples

Key MFA artifacts (planned)

- Considerations around outsourced authentication
- Accessibility evaluation of MFA technologies
- FERPA and MFA contract language
- Funding Models
- Sample project & deployment plans
- Sample support documentation & processes, FAQs, etc.
- Sample user communication campaigns
- Likely to come up with others

MFA Cohortium Webinar

- Webinar last week, which was attended by about 160 people, representing > 100 higher education institutions
- Conducted several polls, lots of interest in the work we are doing

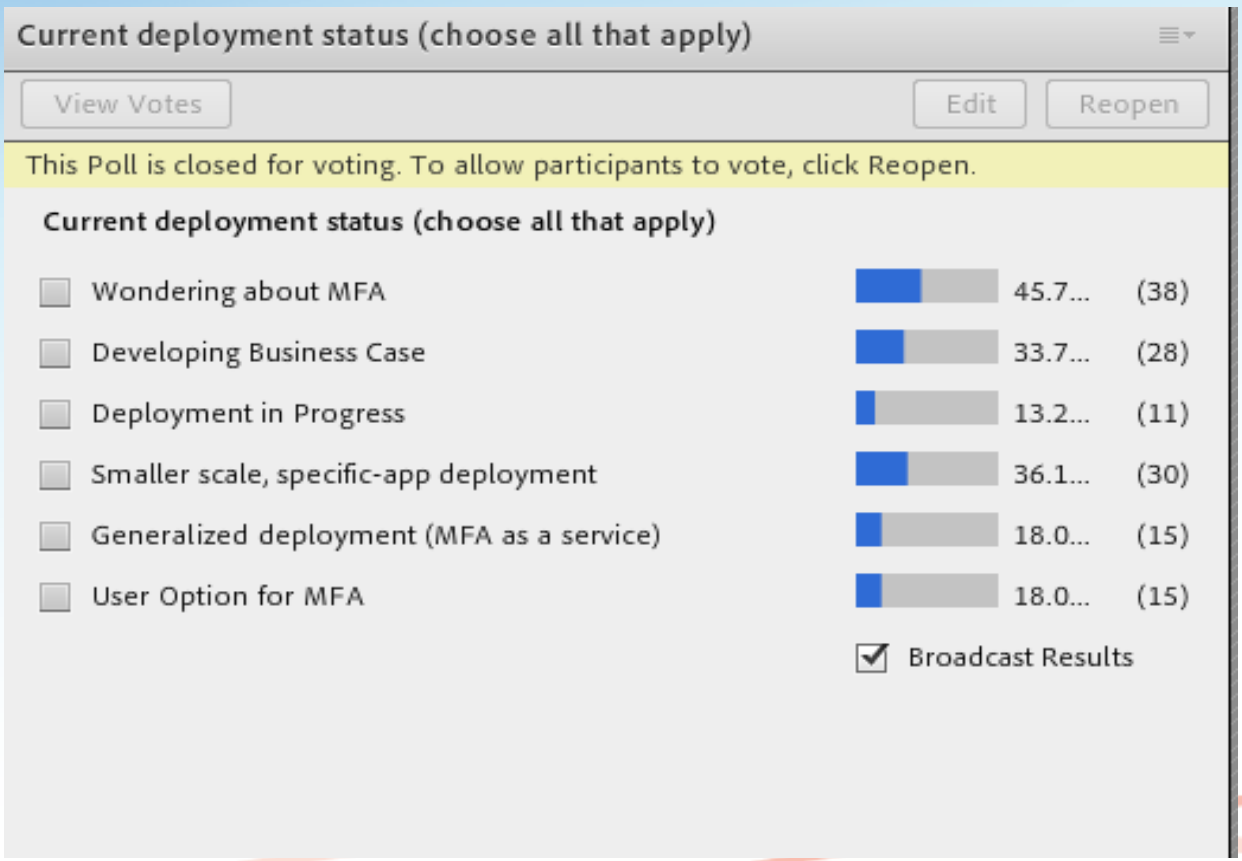

What we hope to learn in the next year

- Compelling business cases
- Effective deployment strategies
  - By risk assessment, by department, by role, by user choice
- Approaches to accessibility
- ROI for various deployment options
- Effective user adoption approaches
- Polished open source ecosystem tools
- Resources which minimize impact on help desk/support infrastructure & any new roles required
- Many things we did not expect…

INTERNET 2

Citizen-centric attribute deliverables

- Schema Catalog and Attribute Registry
  https://spaces.internet2.edu/x/dgROAg
- Attribute-annotated Use Cases
- Cookbook “To Serve Citizens” ☺
- Global Public Inclusive Infrastructure (GPII) Proof of concept, using User-Managed Access (UMA)
- Bindings and refactoring
- Engagement with the privacy manager

Categories of use cases

- Accessibility
  - Physical, cognitive, age-related, etc.
  - Global Publically Inclusive Internet (gpii.net)
- Operational Government
  - Transaction based
  - May be out of scope
- "Social Government"
  - Community wikis, on-line discussions, news feeds, etc
  - Generally local in nature, often requiring anonymous but attribute-controlled access (e.g. resident, registered voter, over legal age, etc.)
- Envision It Scenarios
  - Contained in Full NSTIC Strategy (April 2011)
  - http://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf
- UMA developed
- IdeSG provided – still under distillation

INTERNET 2

GPII Proof of Concept

- The purpose of the Global Public Inclusive Infrastructure (GPII) is to ensure that everyone who faces accessibility barriers due to ***disability***, ***literacy***, ***digital literacy,*** or ***aging***, regardless of ***economic resources***, can access and use the Internet and all its information, communities, and services for education, employment, daily living, civic participation, health, and safety
- Automatic personalization of user interfaces and user context adaptation based on user preferences, across platforms
- Schema standard is AccessForAll (ISO/IEC JTC1 24751)
- http://gpii.net
- Pilot applications, proofs of concept beginning with
  - User preferences stored and accessed securely in an online repository
  - Those preferences drive presentation features that provide accessibility accommodations when user visits online resources
  - All leveraging UMA profiles of Oauth 2.0 aligned with emerging GPII security and privacy architectures

INTERNET 2

Important citizen-centric resources

- The citizen-centric attribute registry
  - http://macedir.org/ontologies/attribute/2012-11-10/attributeOntologyDoc/
- Annotated use cases
  - Intended to capture the issues around the value of various attributes and identifiers in addressing use case issues
  - Several perspectives: user, IdP, SP
  - Real world engineered, from sources of authority to current practices and incentives
  - Could catalyze convergence conversations
  - https://spaces.internet2.edu/display/scalepriv/Annotated+Use+Cases
- Cookbooks and recipes to come
- GPII.net and other resources

What we hope to learn in the next year

- Annotate additional use cases
- Foster some convergence discussions
- Develop key data-driven issues:
  - In R&E, IdP's are normalized (syntax and semantics) on key attributes but differ among themselves in privacy policies and what we release to others; in the social space, IdP's are wildly divergent on attributes but generally promiscuous in which attributes are released.
  - Is there a hierarchical "sweet spot" where users can actively manage privacy with almost no impedance?
  - Internationalization issues, from policy to the Spanish surname topic
- Foster active research on usability within the academic community
- The relationship of citizen-centric attributes to provisioning data

Privacy managers (Carnegie-Mellon University)

- Consoles to help users manage the release of attributes
- Can leverage trust, informed consent, default settings and preferences, etc.
- Must be carefully engineered
  - Across the variety of contexts
  - Across a variety of credential types
  - In ways that are user-effective
- Similar, less leveraged approaches are successfully deployed in a few settings, demonstrating that users can and will manage privacy.
- Research shows that over 90% of social network users do not know what is being released or how to change it

Key design considerations

- Usability
- CMU Tech Report, Warning Design Guidelines, Bauer et al
- Informed and * consent
- GPII
- Technology agnostic – SAML, anon creds, OpenId, etc., though plumbed to Shib to start
- Awareness of out-of-band considerations
- “Nudging” applied to privacy
- Minimal disclosure for constrained purpose
- First alpha due this month

Metadata and trust implications

- At scale, there needs to be ways to establish and convey trusted information about applications and services to users
  - Implies "vetting" or auditing processes for services
  - Implies metadata that can convey this information in real time to users
  - Implies trust in the metadata
- Dynamic metadata services
  - Work is already underway on this in other places
- Federation operations need to evolve
- Auditing applications
  - For "privacy-preserving" approaches (minimal attribute requests, informed consent, proper handling and disposal, etc.), for COPPA compliance, for …
  - Prototype approaches are successful; market needs to grow

INTERNET 2

Requiring care and urgency

- Taking care
  - Federated identity has become core infrastructure; modifying it requires care
  - Integrating with the Shib IdP v3 depends on timing of release (2014) and integration (likely done in parallel)
  - How much can a user intelligently manage? How much do they want to?
  - Getting the defaults right
- Requiring urgency
  - The pain of attribute release is exponentially increasing
  - Consent is the only recourse, as ordained by EU privacy law, campus duchy policy, enterprise lawyer rule, etc.

What we hope to learn in the next year

- Active testing and successive refinements of the privacy manager
  - Nudging approaches
    - By metadata, reputation systems,
  - Able to span multiple technologies, including OpenId Connect
  - Accessibility issues
  - Usability issues, including design, labels, defaults, etc
- Integration with UMA-oriented approaches
- Significant pilot deployments
- Documented interfaces for developing alt privacy managers

Anonymity, unlinkability, and unobservability

- *Anonymity* assures that public data cannot be related to the owner.
- *Unlinkability* assures that two or more related events in an information processing system cannot be related to each other.
  - *Untraceability* assures that two or more events at autonomous systems by the same user cannot be correlated
- *Unobservability* assures that an observer is unable to identify or infer the identities of the parties involved in a transaction.

INTERNET 2

Anonymous Credentials

- Special credentials issued by attribute authorities
- Allows for minimum disclosure of attributes of bearer
  - Over legal age; graduate of university in year X; resident; first-responder certifications; access to age-restricted services; etc
- Can develop trusted responses to access policy by processing previously obtained credentials
  - Eg. Age > 21 developed from birthdate
  - Can use multiple credentials as input when responding
  - Responses optionally contain original attribute values
- Built on several similar technologies, including ABC4Trust (funded by the EU) and uProve (open licensed from MS)
- Tamper-proof
- Unobservable
- Long-time cool technology in search of use cases and modern enhancements (mobility, informed consent, etc.)
- Several pilots looking at integrating them in various ways
- Our work is being led by Brown University

INTERNET 2

Deployment Models

- Classic ABC4Trust, Idemix, etc.
  - Credentials held in a cert store on the user's desktop or smart card
  - RPs accessed via Web Browser
  - Processing done in User's desktop by previously downloaded plugins
- Enterprise-based
  - Credentials held in enterprise directory
  - Processing still done in desktop
  - Addresses mobility
  - May serve important enterprise needs
- Cloud-based
  - Processing and storage moved to the cloud
  - Addresses mobility issues, new devices
- Card based
  - Some way cool smartcard based Dutch work
  - http://www.irmacard.org

INTERNET 2

The Dutch national card

- Technology
  - Smart card with only a photo visible
  - Anonymous credentials from a number of trusted sources
  - Trust roots on card; off-line use
  - Cards interact with apps on Android phone readers via NFC
  - Cards introduce other physical threat vectors
  - Irmacard.org
- Policy and funding
  - The final RFP is being evaluated now; report by end of year
  - Irmacards one of the finalists
  - There are other national cards (e.g. transport) that come in both identity and anonymous flavors.

What we've learned about anon creds

- Badly misnamed technology
  - Can provide identity info, with user consent
  - Provides for minimal disclosure of attributes
  - Lots of alt approaches that use similar phrases such as zero-knowledge, anonymous credentials, double blind gateways, etc.
- The open source is not ready for prime time; the proprietary implementations have lots of issues
- Adding modern features such as mobility and * consent affects trust issues and are poorly addressed
- Deployment model influences trust model
- Still appear to be the best answer for *unobservability*
- Abc4Trust has Inspector mechanism, under user control, allowing for "opening" a policy response

What we hope to learn in the next year

- Expand the use cases to illustrate the potential and the remaining barriers
- Get a reliable and robust open source enterprise-centric platform established
- Understand and evaluate the “privacy leakage” implications of informed consent
- Gain a better understanding of the tradeoffs between desktop and cloud-based processing

Takeaways

- Moving the needle on MFA
  - A number of important, solvable issues are emerging
- Attributes are the key and its already a mess
- Researching what it takes to put the “informed” into consent, and trying to build it
- Anonymous credentials are still immature, and still the only answer to unobservability
- New businesses, such as application auditing, are needed
- The real steady state future is “interfederated identity” but getting there is getting harder