

Agenda – Day One

09:00 – 11:00 Bringing IAM and SSO to the Campus

- 09:00 Welcome and Introductions
- 09:25 Policies at the Local Level
- 10:00 Technologies at the Local Level
- *10:45-11:00 Tea Break*

11:00 – 14:00 Applied Knowledge: Cloud Services

- 11:00 Cloud Services and SSO
- *12:00-13:30 Lunch*
- 13:30 Campuses and the Research Community

14:00 – 16:00 Best Practices and Business Plans

- 14:00 Best Practices and Joining a Federation
- *14:45-15:00 Tea Break*
- 15:00 Business Planning

16:00 – 17:00 Open Discussion

THE CAMPUS IDENTITY SYSTEM

Lucy Lynch, NSRC

Heather Flanagan, NSRC & REFEDS



REFEDS

Learning Objectives



Discovering the
key role campus
networks play in
trusted identities
for R&E

Authoritative source
for user data
associated with your
domain

Logical
implementers and
maintainers of
Identity Management
related core services

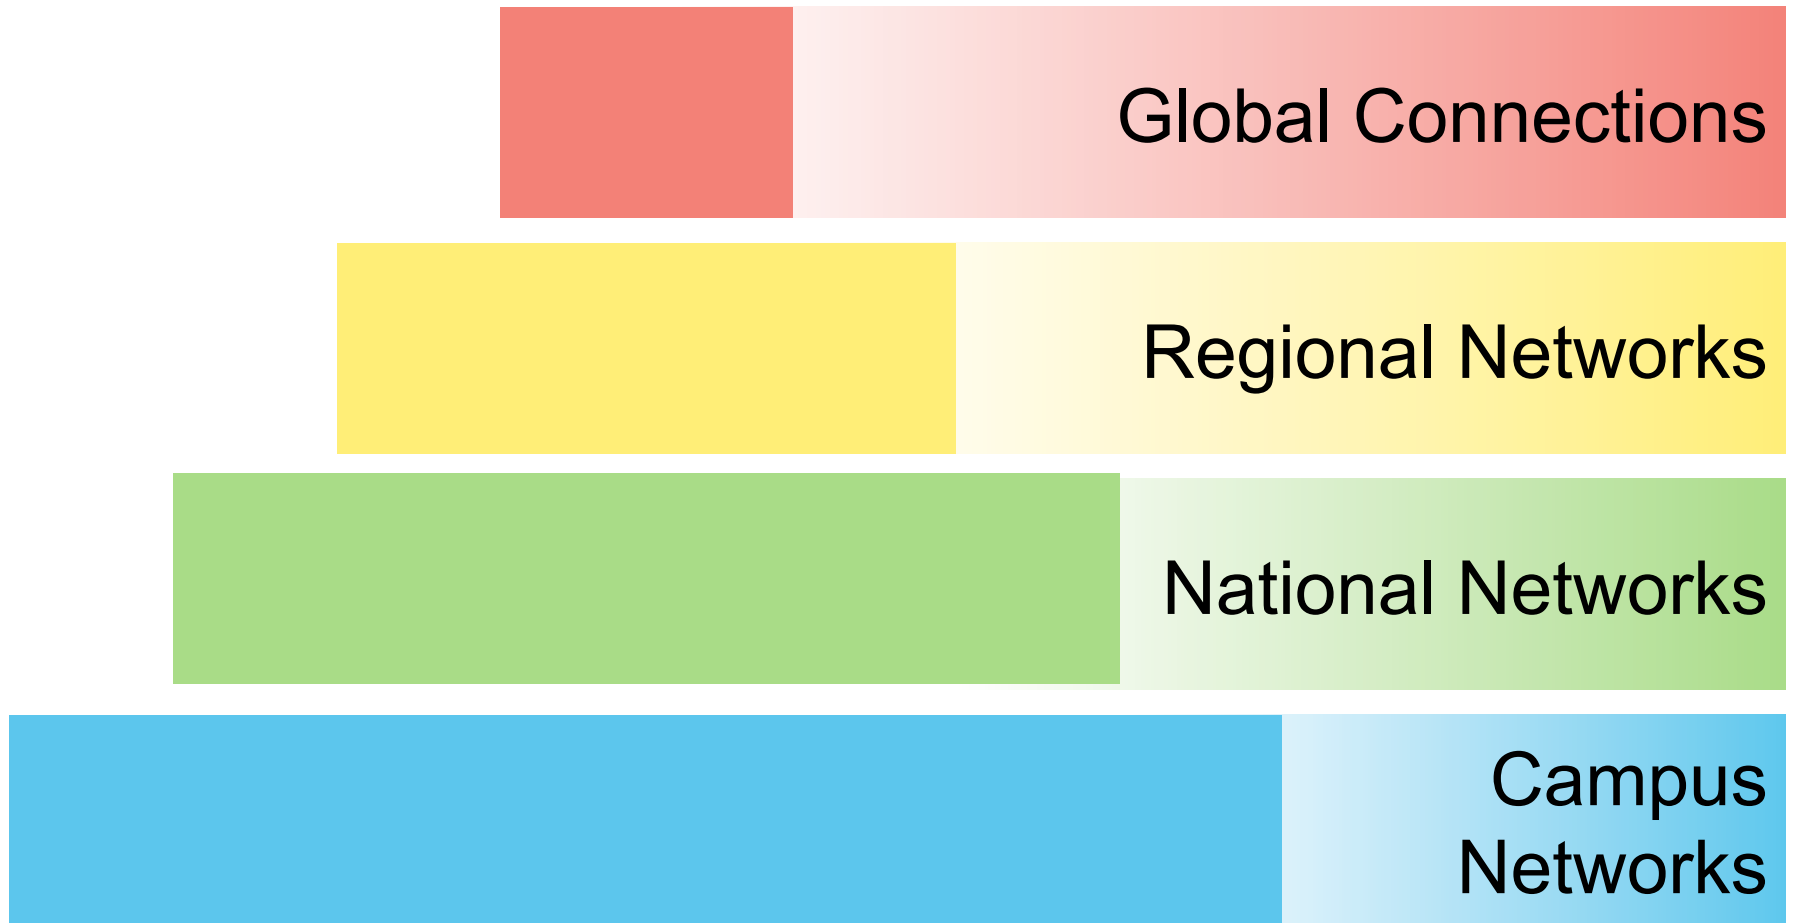
Logical trusted
partners for data
exchange in the R&E
network hierarchy

Research and Education Networking

Global R&E collaborations are based on a bottom up model that manages connections and services in order to provide efficiency and scale using a layered model including:

- Regional RENS
- National Research and Education Networks
- Users at the campus network level
- Global Connectivity

The NREN Ecosystem



Design Goals for Identity and R&E

The dream - providing users with a single login that grants access to any resource, irrespective of device or physical location.

When designing for Identity Management (IdM) start with your desired end goals and then work backwards.

- Single Sign On (SSO)
- Role-based access to network resources
- Support for traveling scholars (think “eduroam”)
- Tools for collaboration
- Shared access to remote instruments
- *Your wish list goes here*



User: Courtney_Alice

Why Focus on Campus Networks?

- Individual institutions are the authoritative source for domain data
- The campus network is the foundation for research and education activities
- The best path to network capacity, equipment and personnel
- No researcher is connected directly to a national R&E network
 - They are all connected to campus or enterprise networks for access

Benefits for Campus Network Operators

When staff and money are in short supply, any new effort must add value to entire campus plan. IdM can provide:

- better utilization data
- better security
- better management for restricted resources

These things come at a cost as there are new services and software to manage and someone will have to maintain data integrity on an on-going basis.

The value goes beyond IdM.

We Already Understand the Model

A good network design is modular and hierarchical, with a clear separation of functions: core, distribution, and access.

Good campus networks will leverage:

- Domain-based span of control
- Layered services built around your core
- Scalable, interoperable, standards-based technical choices

The Identity Management model is much the same

Identity Management Services Capabilities

- Centered on the User Identifier (NetID) - A single unique University wide identifier bound to the individual user and used at log-in to provision:
 - Authentication
 - Quickly verify user identities (Who you are)
 - Authorization
 - Control users access (What you can access)
 - Administration
 - Manage user privileges by role, group, status, etc.
 - Allows for fine-grained policy application

One Way to Think About It -

Interdomain Routing

- . IGP/iBGP
- . ASN
- . eBGP

Identity management

- . Campus IdM
- . Federation
- . Inter-Federation

If you are a network engineer you all ready deal with local policy and global transit as part of your day job.

You would not allow an unmanaged device on your network
– why allow an unmanaged end user?

The Late Mover Advantage

In the last ten years R&E networks have seen a lot of progress in Identity Management.

- Common Standards
- Common Software
 - With Open Source options!
- Common Profiles
- Common Practices and Policies

New entrants benefit from the lessons already learned

Case Examples

As you dig into the details of Identity Management you may be interested in specific examples of both campus and NREN deployments. There are many successful cases to choose from but these two have excellent documentation with robust links to resources.

- NREN: Canadian Access Federation (CAF) – CANARIE
 - <http://canarie.ca/identity/caf/>
 - Also includes links to a packaged solution using common tools
<https://github.com/canariecaf>
- Campus: Rutgers, The State University of New Jersey
 - <https://idms.rutgers.edu/>

Communities of Practice

The R&E community has several well developed forums for Identity practitioners which are open to new participants. These forums include training resources, special advanced topic working groups, and documentation on current best practices. The sites provide both technical and policy guidance.

- REFEDS (Research and Education FEDerations group)
 - EU-based group <https://refeds.org/>
- InCommon(operated by Internet2 Staff)
 - US/Internet2 Based group <https://www.incommon.org/>
- eduroam (education roaming)
 - secure, world-wide roaming access service
 - <https://www.eduroam.org/>
- eduGAIN (operated by GÉANT)
 - interconnects identity federations around the world
 - <http://www.eduGAIN.org/>

The Tool Box

There are many ways to put together Authentication and Authorization services and lots of options for centralized data management. There are also good open source tools for identity federation management. These tools rely on your underlying network and wireless infrastructure and can be customized to match your campus plan. Two commonly used examples:

- Shibboleth: Federated Services (IdP/SP)
 - <https://shibboleth.net/>
- CAS (central authentication service for SSO)
 - <http://www.ja-sig.org/products/cas/overview/index.html>

Building Identity Block by Block

Elements of IdM

- unique identifier
- directory
- authentication
- password store
- authorization
- federation
 - identity provider
 - service provider
 - directory service

Deployment Examples

- netid
- ldap
- cas
- kerberos
- mysql
- shibboleth

Note that all of the examples require customization based on your local policy

Getting Started

Create a campus inventory that includes:

- Your existing data sources
- Your current authentication sites and methods
- Your current authorization policies and methods
- Your existing software and services
- A survey of your users to gather requirements for both internal and external identity based access
- Your institutional policies on user data
 - Including privacy, security, and acceptable use

TECHNOLOGY ON CAMPUS

(with many thanks to Chris Phillips and Tom Barton, “Demystifying Privilege and Access Management” IAM Online session in August 2012)

Learning Objectives

Understanding
how Campus
Identity Systems
evolve

Campus IAM

- Identities
 - Registries of who/what, identifiers, attributes, systems integration
- Credentials & authentication
 - Internal, external
 - Linked to Identity
- Access management
 - Roles, rules, entitlement, affiliation, groups, privileges, policy, authority, delegation, etc.

Factors to Consider in IAM Platforms

- Variable distance from the data
 - Local - Systems of Record, home grown, commercial
 - Internally, apps enjoy tightly coupled access to fresh data
 - Usually 'behind the wall'
 - Federated – apps aware of more than local & remote systems of record
 - accept other identities underpinned by a trust decision
 - Foot in both worlds at times -- inside & outside the wall.
 - Data 'distance' further out
 - Upon sign-in of user, or provisioning task
 - Cloud (aka SaaS, PaaS) - apps abstracted away lower level details, could be furthest away from your fresh data
 - Similar challenges as federated
 - SLAs may or may not be under your control
 - Outside the wall
 - Deployment profile & app sophistication guide data management

Factors: Design and Intent

- Different design patterns
 - Rigid data structures vs. elastic ones
 - Sometimes flexibility to a fault – not prescriptive enough
- Philosophical design differences
 - Intentionally designed to support externalized AuthN/Z vs. bolt on
- Implementer intentions
 - Walled garden product may cause challenges
- Effort to keep current
 - Balance between get it done now vs. perfect design

Factors: Governance and Process

- Governance
 - Clarity around:
 - Who says who says
 - Authorization model:
 - Centralized or distributed?
 - Application or data centric?
- Process and Practices
 - Are change control practices in place & recognize the implications of local, federated, and cloud use styles?
 - System Of Record steward may not realize:
 - Dependencies on their data and change turnaround
 - How far flung systems of record data may be

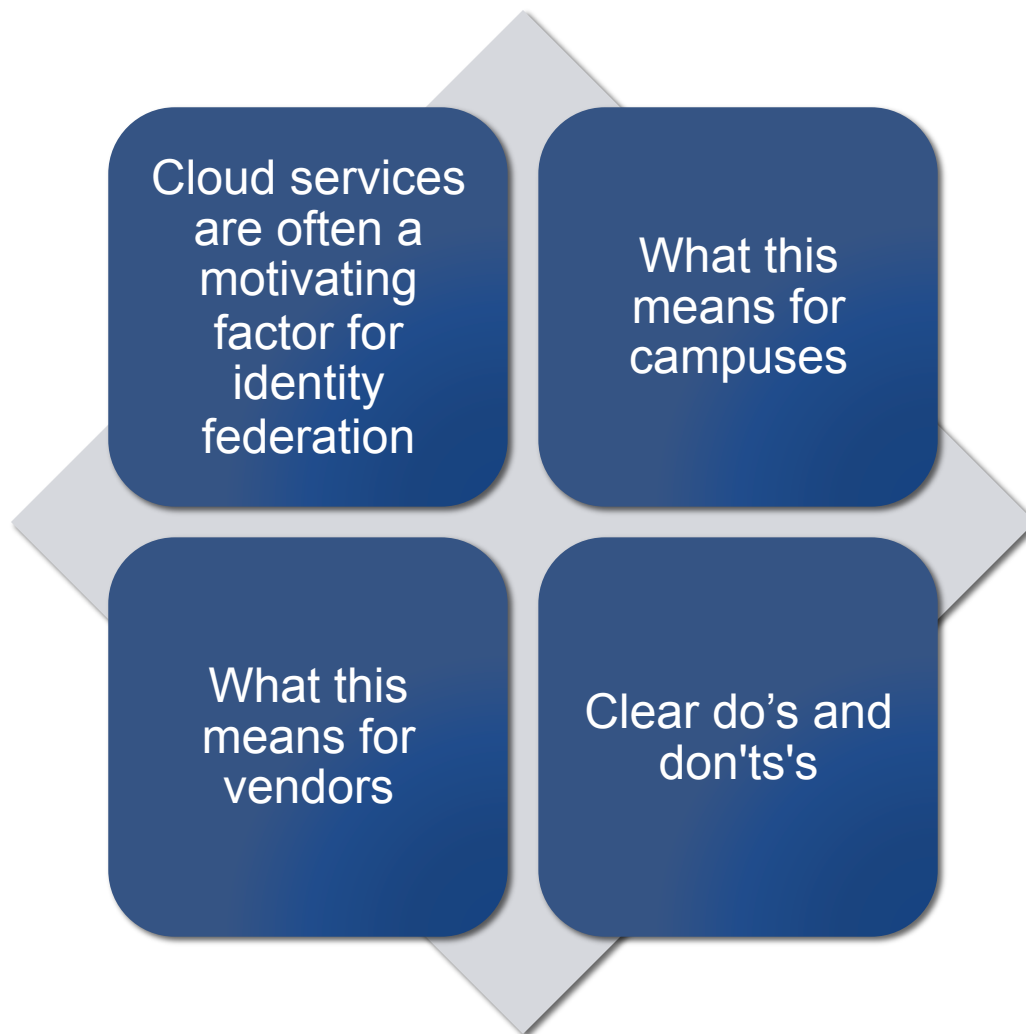
The Evolution of Access Management

Phase	Description
None – most physical controls	If you can authenticate, you get everything
Control by contract	If you can authenticate, you get everything, but there is no abuse policy in place
Hard coded privilege tables at the resource	Authorization at the application level
The above + LDAP calls for intrinsic attributes	Authorization starts to depend on external attributes
An attribute authority	An application or service can get any attribute that the access management policies permit
An external yes/no authorization service	An external service calculates whether access is permitted

IDENTITY AND CLOUD SERVICES

<https://wiki.refeds.org/display/FBP/Cloud+Services+Cookbook>

Learning Objectives



Cloud Services – Easy?

- Not generally – can be complicated and time consuming
- Provisioning individually per service is difficult on both the user and the campus
- Using a shared account let's the service provider pick their favorite number around licensing
- Think about control over provisioning, de-provisioning, and support for the privacy requirements the campus has to follow

Handling Authentication

DON'T assume successful authentication means the user is authorized for service.

- Campus: CONSIDER stronger authentication (e.g., multi-factor) over password strengthening (increasing length, complexity requirements)
- Vendor: DO let the identity provider handle authentication
- Vendor: DO rely on browser-based authentication for non-browser applications.
- Vendor: DON'T use service-specific passwords unless there are no alternatives.
- Vendor: DO use forced re-authentication when appropriate.

Identifiers

- Both Campus and Vendor: DO support a varied set of identifiers.
- Both Campus and Vendor: CONSIDER the use of eduPersonTargetedID where appropriate.
- Both Campus and Vendor: DO use standard definitions of identifiers and attributes.
- Both Campus and Vendor: DON'T mistake eduPersonPrincipalName for a valid email address.
- Campus: DO standardize internally on a stable "serial number" for users.
- Campus: DO make eduPersonPrincipalName useful.

Authorization

Remember: AuthZ is not the same as AuthN!

- Both Campus and Vendor: DO leverage eduPerson attributes for authorization.
- Both Campus and Vendor: DO be clear about where the allow/deny decision logic is evaluated.
- Both Campus and Vendor: DO determine whether a service is dependent on service-specific "local" user accounts.

Provisioning and De-provisioning

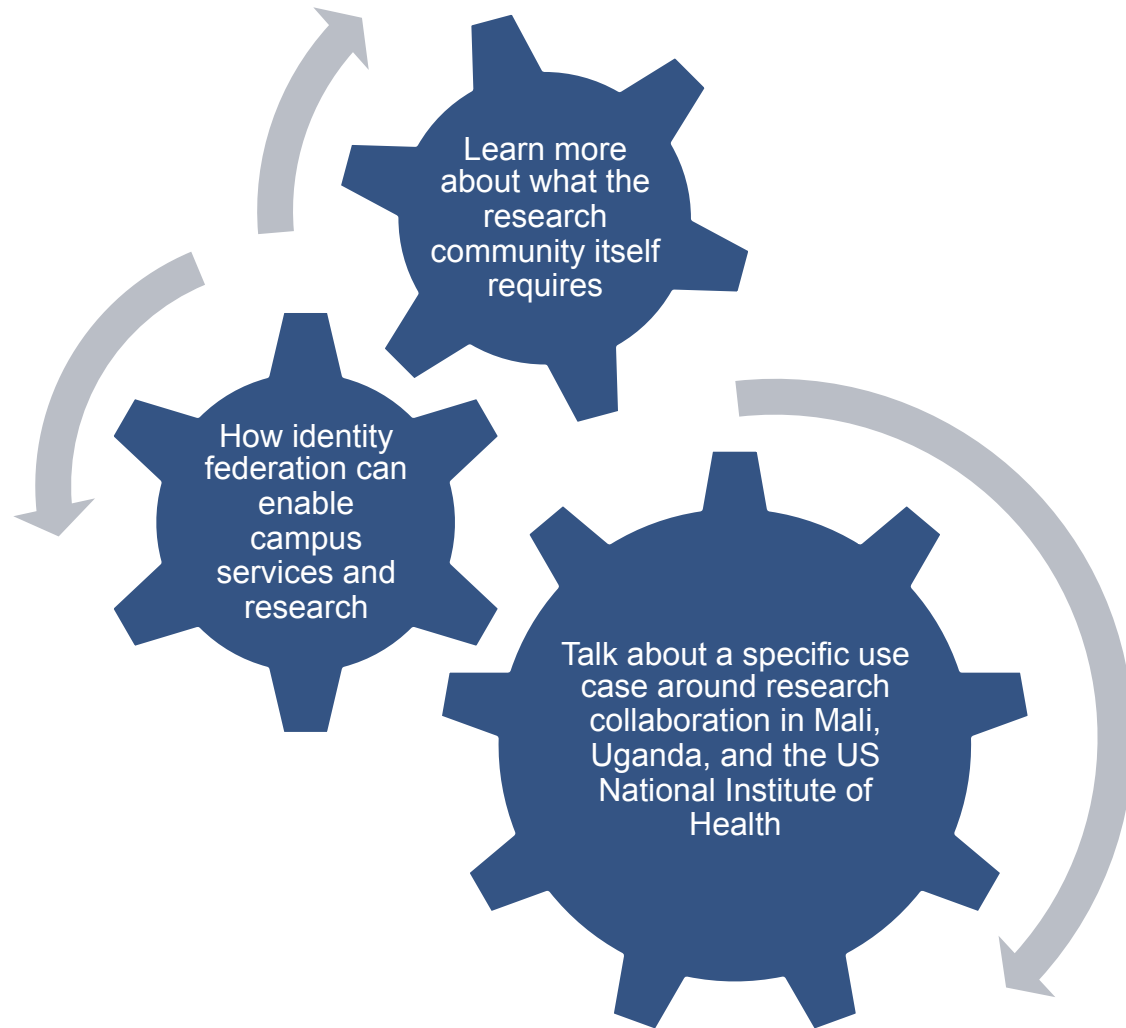
- Campus: DO expect the typical vendor to have a single, set model for creating user accounts on their systems.
- Campus: DO practice "defensive programming" when setting up provisioning services.
- Campus: DON'T require out-of-band acceptance of Terms of Use.
- Campus: DON'T expect robust de-provisioning support.
- Campus: DO handle username changes.
- Vendor: DO support just-in-time provisioning based on user attributes passed in SAML assertions whenever possible.
- Vendor: DO consider standardizing your provisioning (and de-provisioning) APIs.
- Vendor: DO manage your provisioning API in a way that respects the service subscriber interests.

RESEARCH AND COLLABORATIONS

Heather Flanagan, NSRC
Chris Whalen, NIH



Learning Objectives



Virtual Organization and Research Groups

- The “Buried Scholar” problem
 - Improve the reputation of both the campus and the researcher by having a campus-branded identity
 - Some regions start with having a regionally branded identity, pooling the resources to focus on SSO first, and site-specific support later
 - More efficiently allocate resources by having federated identity as a campus or NREN-based service, rather than having individual research departments build this on their own
- Global science and research opportunities
 - Scientists can more easily participate in global collaboration when the collaborations do not have to set up individual accounts

Research on Researcher Needs

- Original FIM4R paper in 2012 described a set of recommendations to the research communities, technology providers, and funding agencies
 - The core use cases came from large research organizations with funding

<https://cdsweb.cern.ch/record/1442597>

- The “Advancing Technologies and Federated Communities”, also in 2012, described a set of recommendations around technology, policy, funding, and legal issues.
 - A more generalized approach than the FIM paper, but the recommendations are largely the same

<https://www.terena.org/publications/files/2012-AAA-Study-report-final.pdf>

FIM4R Findings Summarized

- Federated technologies are good. Take advantage of them.
- The infrastructure needs to be improved to take advantage of federated technologies. Do it.
- Relying on the older models of local account creation and IP-based ACLs is easier. This is a very limited view. Stop it.
- If you can't fix it all yourself (and you can't), facilitate the efforts of groups that can. Build relationships, target your spending or funding to make the biggest impact.

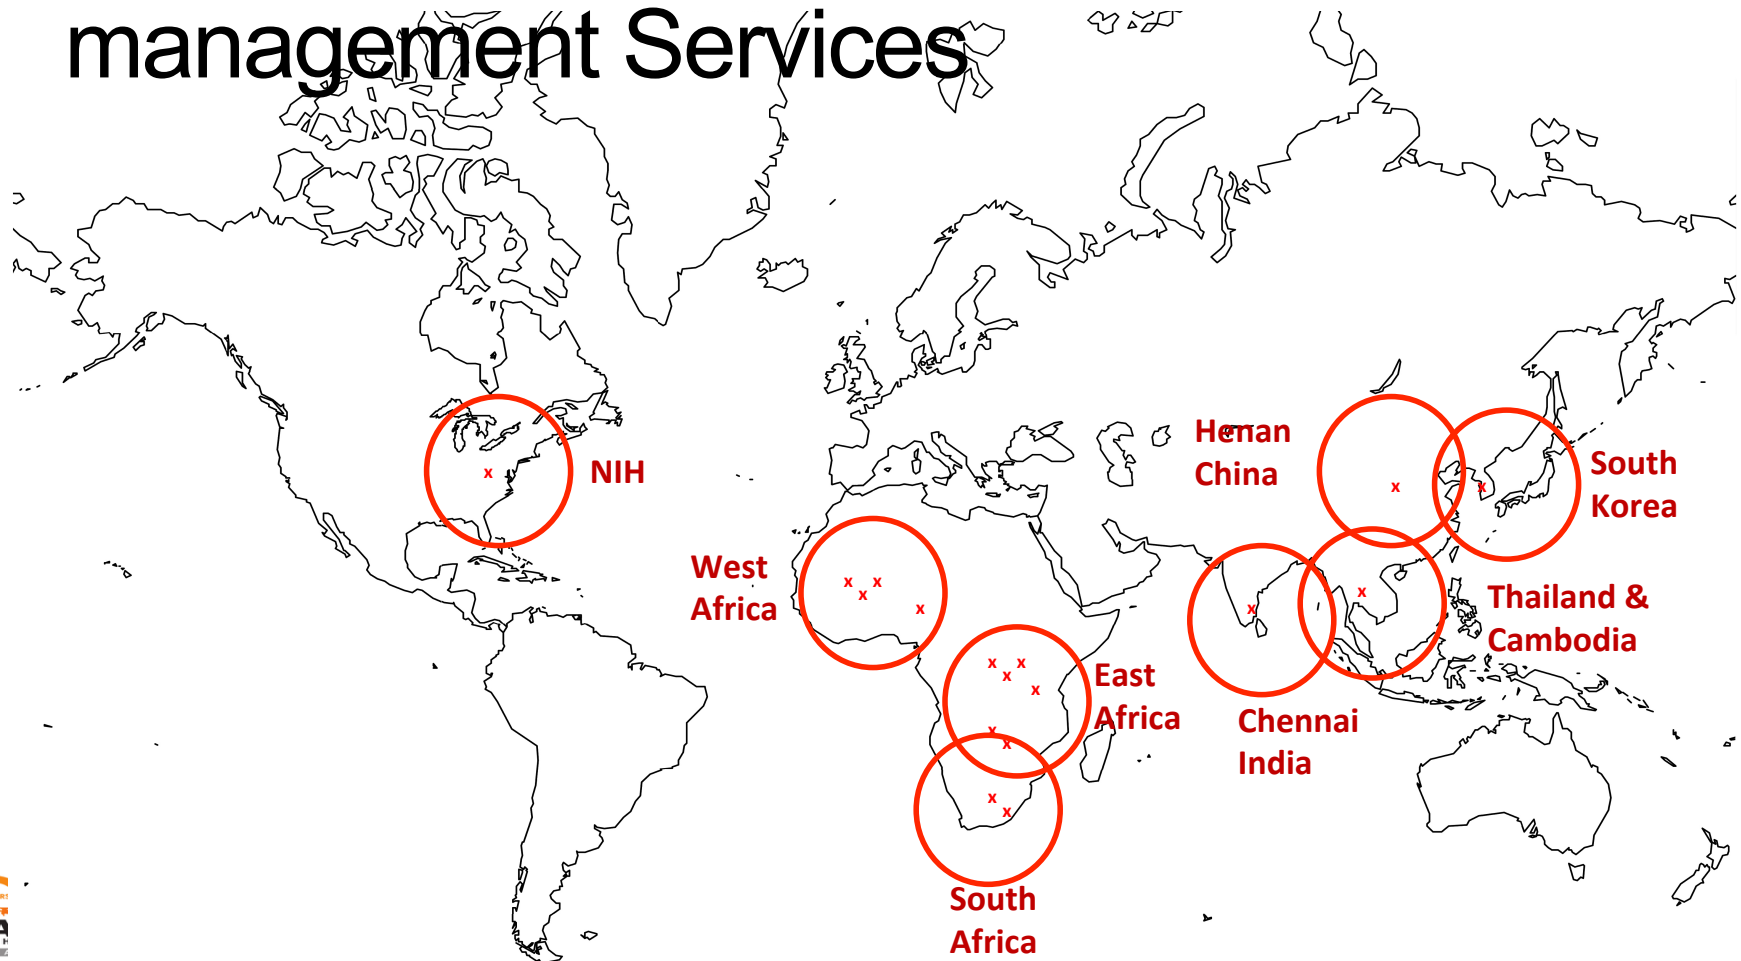
Case Study

- Center for Infectious Diseases (US NIH) and Centers of Excellence in Uganda and Mali
 - Similar work is going on in India and China

Virtual Organizations for collaboration at a West African Research Center

Christopher Whalen
Program Manager
International Biomedical Research Support
Program
NIH/NIAID Office of Cyberinfrastructure and
Computational Biology

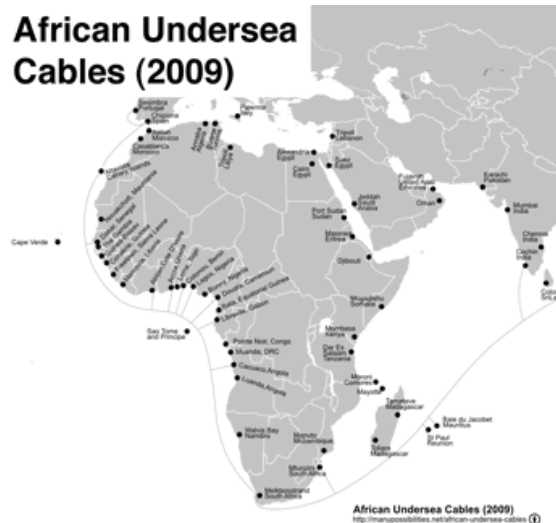
NIAID International Biomedical Research Support Program – Infrastructure and Clinical Data management Services



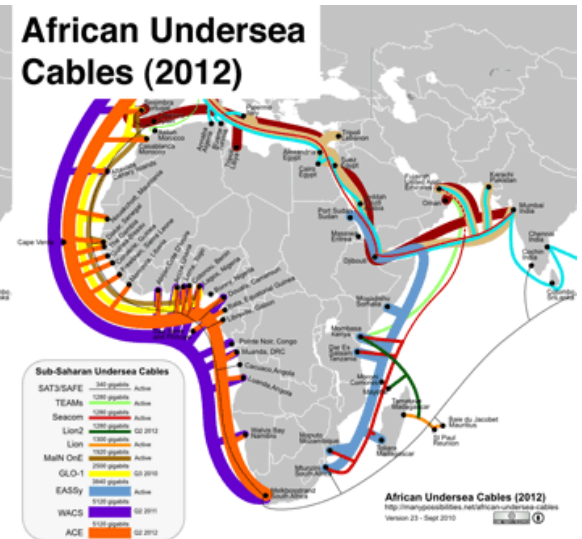
Collaborative networks starting with basic connectivity



African Undersea Cables (2009)



African Undersea Cables (2012)



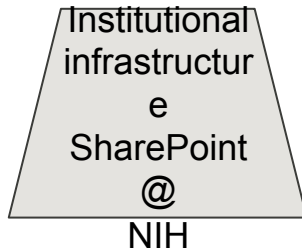
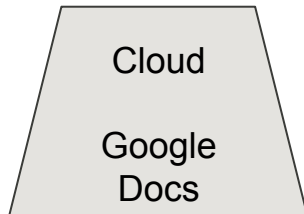
Science depends on collaboration between researchers and institutions

- Collaboration is critical because of the increasing multidisciplinary nature of modern science
 - Example: Molecular Biology research often encompasses both molecular and cell biology including structural and functional genomics, transcriptomics, proteomics, bioinformatics, biomedicine, molecular enzymology, molecular virology and molecular immunology, theoretical bases of biotechnology, physics and physical chemistry of proteins and nucleic acids.

Collaboration Tools for science

- Document Management
 - Storage and retrieval
 - Version Control
 - Search
 - Categorize
- Information Management
 - Calendar
 - Contacts
 - Blogs
 - Discussions
 - Wikis
- Collaboration
 - Access Controls
 - project workspaces
 - Task lists
 - workflows
 - data sharing
 - project discussions

Traditional approach to collaboration tools



Collaboration Identity (username/password)	Institutional Identity (username/password)
cwhalen1234@gmail.com collaborator1234@gmail.com	christopher.whalen@nih.gov collaborator1234@university.ed.uk
christopher.whalen@nih.gov collaborator1234@nihext.nih.gov	christopher.whalen@nih.gov collaborator1234@university.ed.uk

Virtual organization/federation approach to collaboration tools

	Collaboration Identity (username/password)	Institutional Identity (username/password)
Cloud with VO Collaboration Tool	christopher.whalen@nih.gov collaborator1234@university.ed.uk	christopher.whalen@nih.gov collaborator1234@university.ed.uk
Institutional infrastructure SharePoint @ NIH using the VO	christopher.whalen@nih.gov collaborator1234@university.ed.uk	christopher.whalen@nih.gov collaborator1234@university.ed.uk

Scientific Tools for researchers

- Translational Science has experienced an explosion of data
 - Challenges of sharing databases
- High Performance Computing resources
 - Sequencing Analysis
- Data Management tools for clinical trials and studies
- Specimen, compound, databases
- Support Validation and auditing

Virtual Organization Executing Bioinformatics Analysis

Without Virtual Organization (Currently)

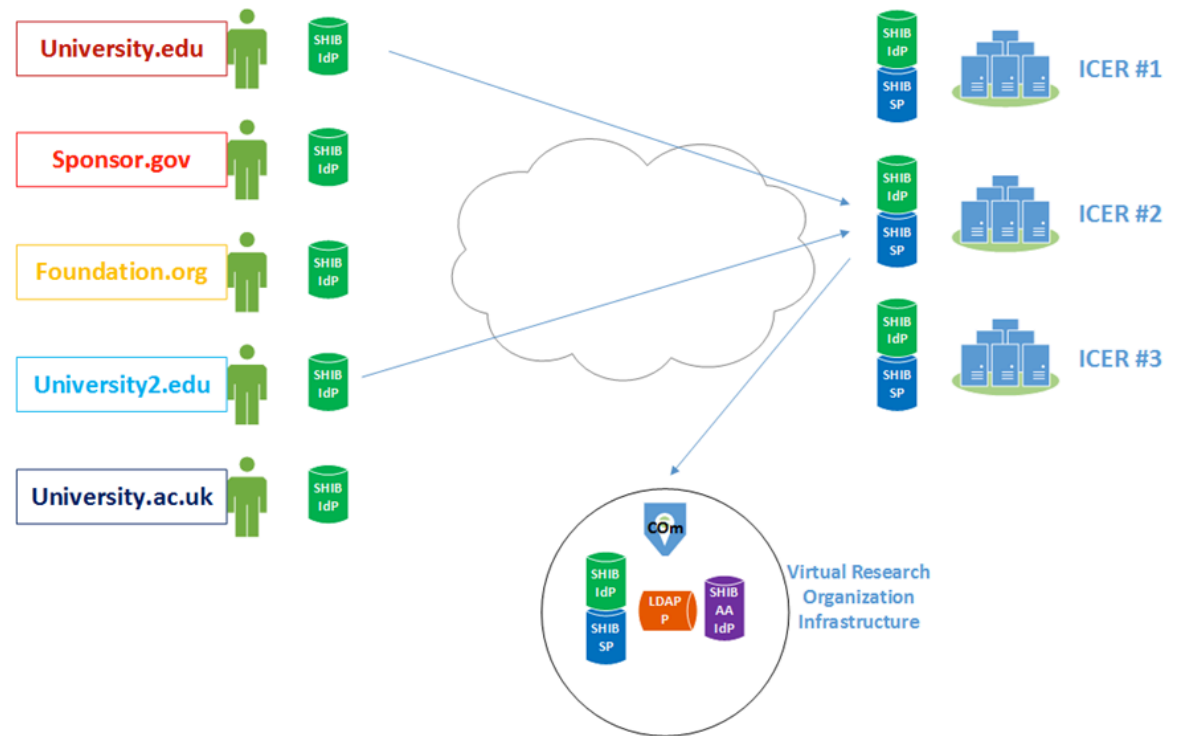
- NIAID HPC login
–ssh squires@hpc.niaid.nih.gov
- Mali ICER login
–ssh icermali/squiresrb@slipstream.icermali.org
- Uganda ICER login
–ssh iceruganda/
squiresrb@slipstream.iceruganda.org

With Virtual Organization

- Log in
–ssh squires@hpc.niaid.nih.gov
- Mali ICER login
–ssh squires@hpc.niaid.nih.gov
- Uganda ICER login
–ssh squires@hpc.niaid.nih.gov

The International Center for Excellence in Research as a Virtual Organization

Modern research collaborations stretch across many different organizations and national boundaries.



The world

Identity Providers IdP



Mali ICER

Uganda ICER

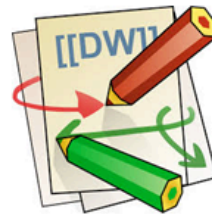


LONDON
SCHOOL of
HYGIENE
& TROPICAL
MEDICINE



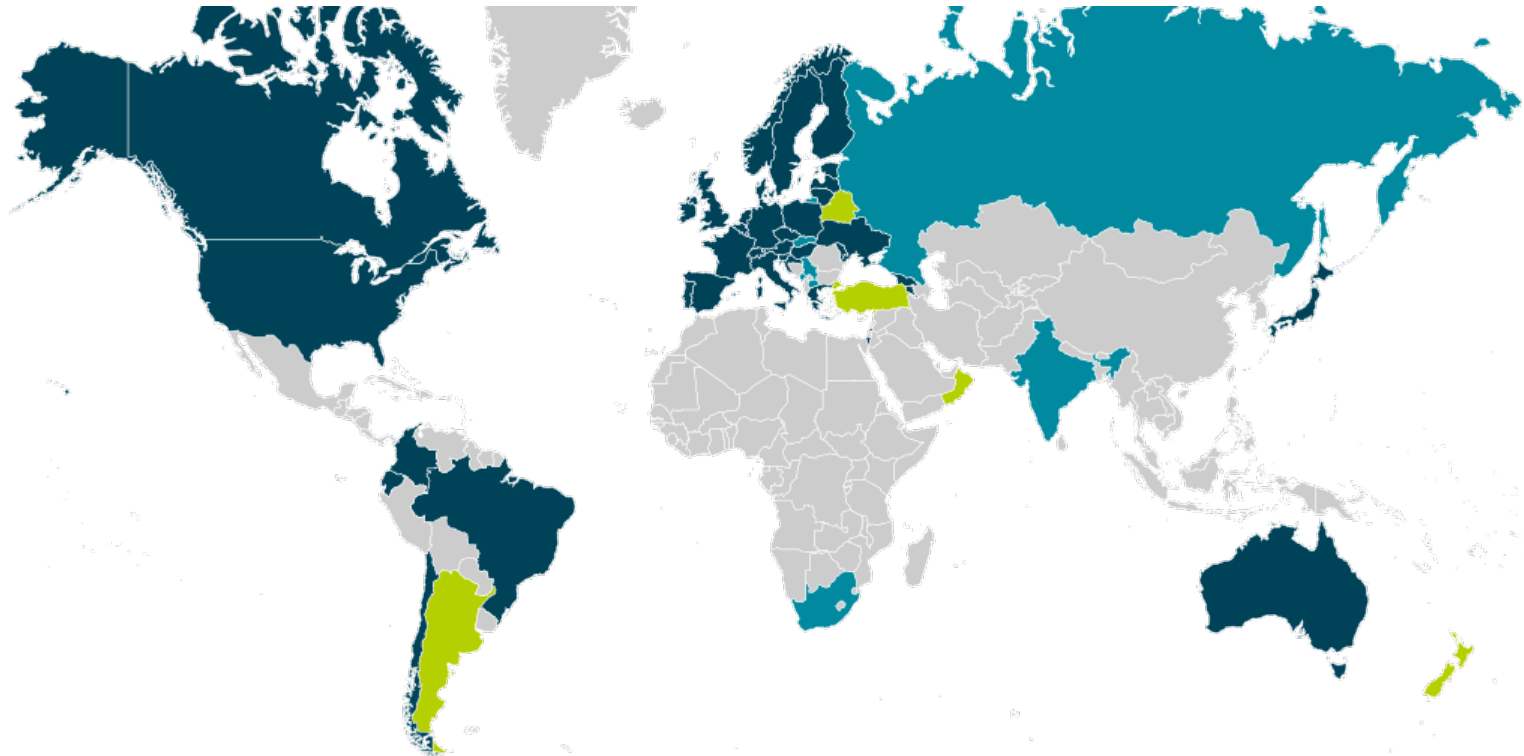
UCSF

Service Providers SP



QlikView

Interfederation: eduGAIN



eduGAIN



Voting-
only



Candidate

Example: Collaboration => Virtual Organization

Polymorphisms in the K13-Propeller Gene in Artemisinin-Susceptible *Plasmodium falciparum* Parasites from Bougoula-Hameau and Bandiagara, Mali

Amed Ouattara, Aminatou Kone, Matthew Adams, Bakary Fofana, Amelia Walling Maiga, Shay Hampton, Drissa Coulibaly, Mahamadou A. Thera, Nouhoum Diallo, Antoine Dara, Issaka Sagara, Jose Pedro Gil, Anders Bjorkman, Shannon Takala-Harrison, Ogobara K. Doumbo, Christopher V. Plowe and Abdoulaye A. Djimde*

Author Affiliations

Department of Epidemiology of Parasitic Diseases, University of Science, Techniques and Technologies of Bamako, Bamako, Mali; Vanderbilt University Medical Center, Nashville, Tennessee; Howard Hughes Medical Institute, Center for Vaccine Development, University of Maryland School of Medicine, Baltimore, Maryland; Department of Microbiology, Tumor and Cell Biology, Karolinska Institutet, Stockholm, Sweden; Department of Physiology and Pharmacology, Karolinska University Hospital / Karolinska Institutet, Stockholm, Sweden; Drug Resistance and Pharmacogenetics, Center for Biodiversity, Functional and Integrative Genomics, Faculdade de Ciências, Universidade de Lisboa, Lisbon, Portugal; The Harpur College of Arts and Sciences, Binghamton University, The State University of New York, Binghamton, New York

Author Notes

* Address correspondence to Abdoulaye A. Djimde, Department of Epidemiology of Parasitic Diseases, University of Science, Techniques and Technologies of Bamako, BP 1805, Bamako, Mali. E-mail: adjimde@icermali.org

University of Science, Techniques, and Technology - Bamako

Abdoulaye A. Djimde
Ogobara K. Doumbo
Mahamadou A. Thera,
Bakary Fofana,
Aminatou Kone,
Amed Ouattara

University of Maryland

Matthew Adams, Shay Hampton, Shannon Takala-Harrison, and Christopher V. Plowe,

Vanderbilt University

Amelia Walling Maiga,

Universidade de Lisboa

Jose Pedro Gil

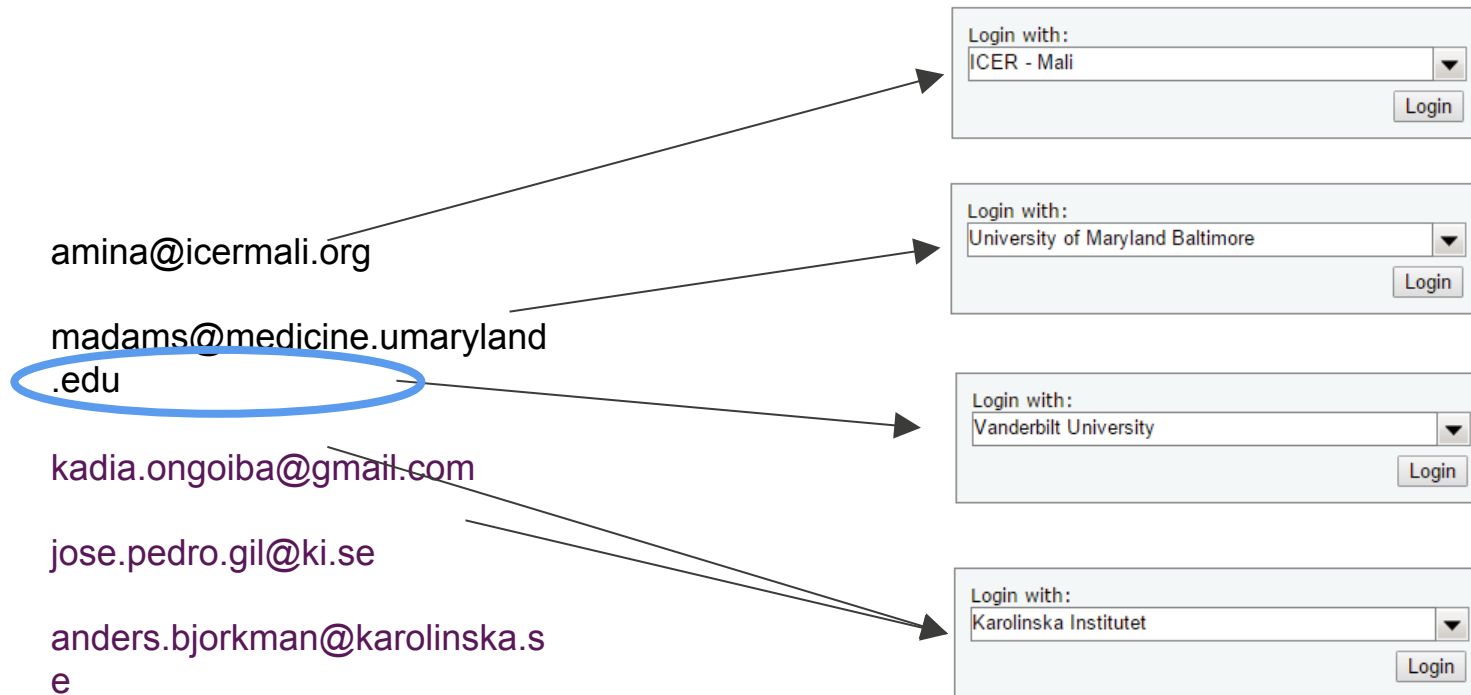
Karolinska Institutet

Anders Bjorkman
Jose Pedro Gil

The State University of New York, Binghamton

Jose Pedro Gil

Virtual Organization – Single username/password



Virtual Organization - Identity consolidation

Universidade de Lisboa

Jose Pedro Gil

The State University of
New York, Binghamton

Jose Pedro Gil

Karolinska Institutet

Jose Pedro Gil

Organizational Identities

Name	Organization	Affiliation	Login Identifiers
Christopher Whalen (23)	National Institute of Allergy and Infectious Diseases		cwhalen@icermali.org (eppn) cwhalen@nih.gov (eppn) CWHALEN@nih.gov (eppn) cwhalen@iceruganda.org (eppn) nihcwhalen@yahoo.com (eppn)

Why NREN Identity?

Why not Google, Facebook, Yahoo!, Twitter..?

Interfederation Research Participants (eduGAIN *via* InCommon, Canarie, etc.) are responsible for investigating and compliance with international privacy law of the countries where research occurs or research subjects reside - *Susan Blair, Chief Privacy Officer, University of Florida, Internet2 Global Summit 2015*

- Google Policy: “Information we collect when you are **signed into Google**, in addition to information we obtain about you from partners, may be associated with your Google Account. When information is associated with your Google Account, we treat it as personal information.”

Acknowledgments

NIAID CIO - Michael Tartakovsky

InCommon - Ann West, Tom Scavo, Ken Klingenstein, John Krienke

Mali ICER IT Group – Sidy Soumare, Amadou Diallo, Mady Diarra,
Salimata Traore, Issa Ba, and Aoua Coulibaly

NIAID OCICB - Matthew Economou, Jake Jester, Jaskiran Singh

Spherical Cow - Heather Flanagan, Scott Koranda, Benn Oshrin

NIH - Jeff Erickson, Sandeep Sathyaprasad

Participants at Advance CAMP 2014 Internet2 Technical Exchange

VOs that have paved the way - LIGO, CERN, Elixir, and others

INTRODUCTION TO IDENTITY FEDERATIONS

Heather Flanagan, NSRC



Learning Objectives

Why identity management is important, and what federations have to offer

How to bring identity management and identity federation to your campus or region

How identity federation can enable campus services and research

How to build a business model that in support of identity federation on your campus

What policies and operational practices you need to have in place

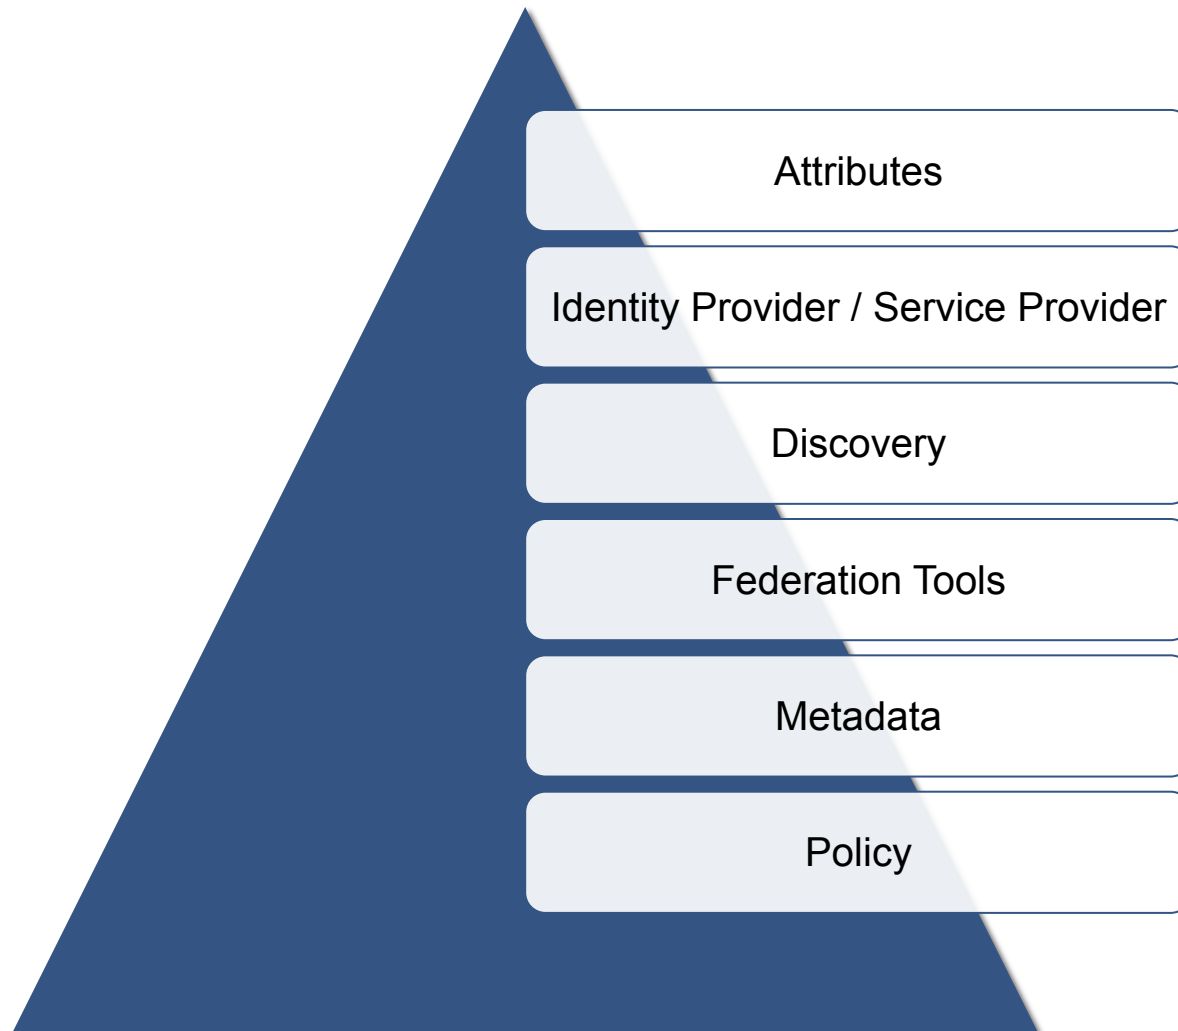
How to engage with the global R&E community

What Is Identity Federation?

“a common framework for trusted shared management of access to on-line resources”
InCommon

“...identities from one organization may use Shibboleth [or another authentication service] to gain federated access to services hosted by another organization. Membership of a federation places obligations on members which *allow* members to trust identity assertions provided by other members.”
JISC

Building Blocks of Federation



Who Benefits? Students and Researchers

Students and researchers

- more collaboration opportunities
- potential access to more resources and data

The research community

- more efficient utilization of resources
- easier research collaboration – can be setup within hours rather than days/weeks
- easier to share or move data between sites/nodes - where relevant

Who Benefits? The Campus

The campus or institution

- a solidly branded institutional identity which improves the overall reputation of the organization
- a stronger security profile for the network
- an ability to logically budget for the network based on actual data (who is on the system, how quickly is it growing, where are the bottlenecks)
- fewer bilateral contracts; more organizations can function under a common framework

“Identity federation participants could spend time establishing operating principles, technology hooks, and agreed-upon data exchange elements with each partner; or they could do it once through the federation and then leverage these common elements for many relationships.” --

InCommon

Benefits/Compelling Reason to Act

Reduces work

- Authentication-related calls to Penn State University's helpdesk dropped by 85% after they installed Shibboleth

Provides current data

- Studies of applications that maintain user data show that the majority of data is out of date. Are you “protecting” your app with stale data?

Insulation from service compromises

- In FIM data is pushed to services as needed. If those services are compromised the attacker can't get everyone's data.

Minimize attack surface area

- Only the IdP needs to be able to contact user data stores. All effort can be focused on securing this one connection instead of one or more connections per service.

What Are Some Compelling Service Possibilities?

- eduroam
- eduGAIN
- digital libraries
- licensed software
 - Learning Management Systems
 - Wikis
- Cloud service providers supporting research and education
 - Researchresearch.com
 - Qualtrics
 - AWS Research Grants

What do Federations do?

At a minimum a federation maintains the list of which IdPs and SPs are in the federation

Most federations also

- Define agreements, rules, and policies
- Provide some user support (documentation, email list, etc.)
- Operate a central discovery service and test infrastructure

Some federations

- Provide self-service tools for managing IdP and SP data (Resource Registry)
- Provide application integration support
- Host or help with outsourced IdPs (IdP in the Cloud, hosted IdP)
- Provide tools for managing "guest" users
- Develop custom tools for the community



How to Make Federated Identity Work

- Start with establishing campus identity systems
- Base-level requirements:
 - centralized campus or institution identity store (e.g., database, LDAP directory)
 - documented policies regarding the life cycle of organizational identity
 - a business model for ongoing development and support

Additional Reading Material

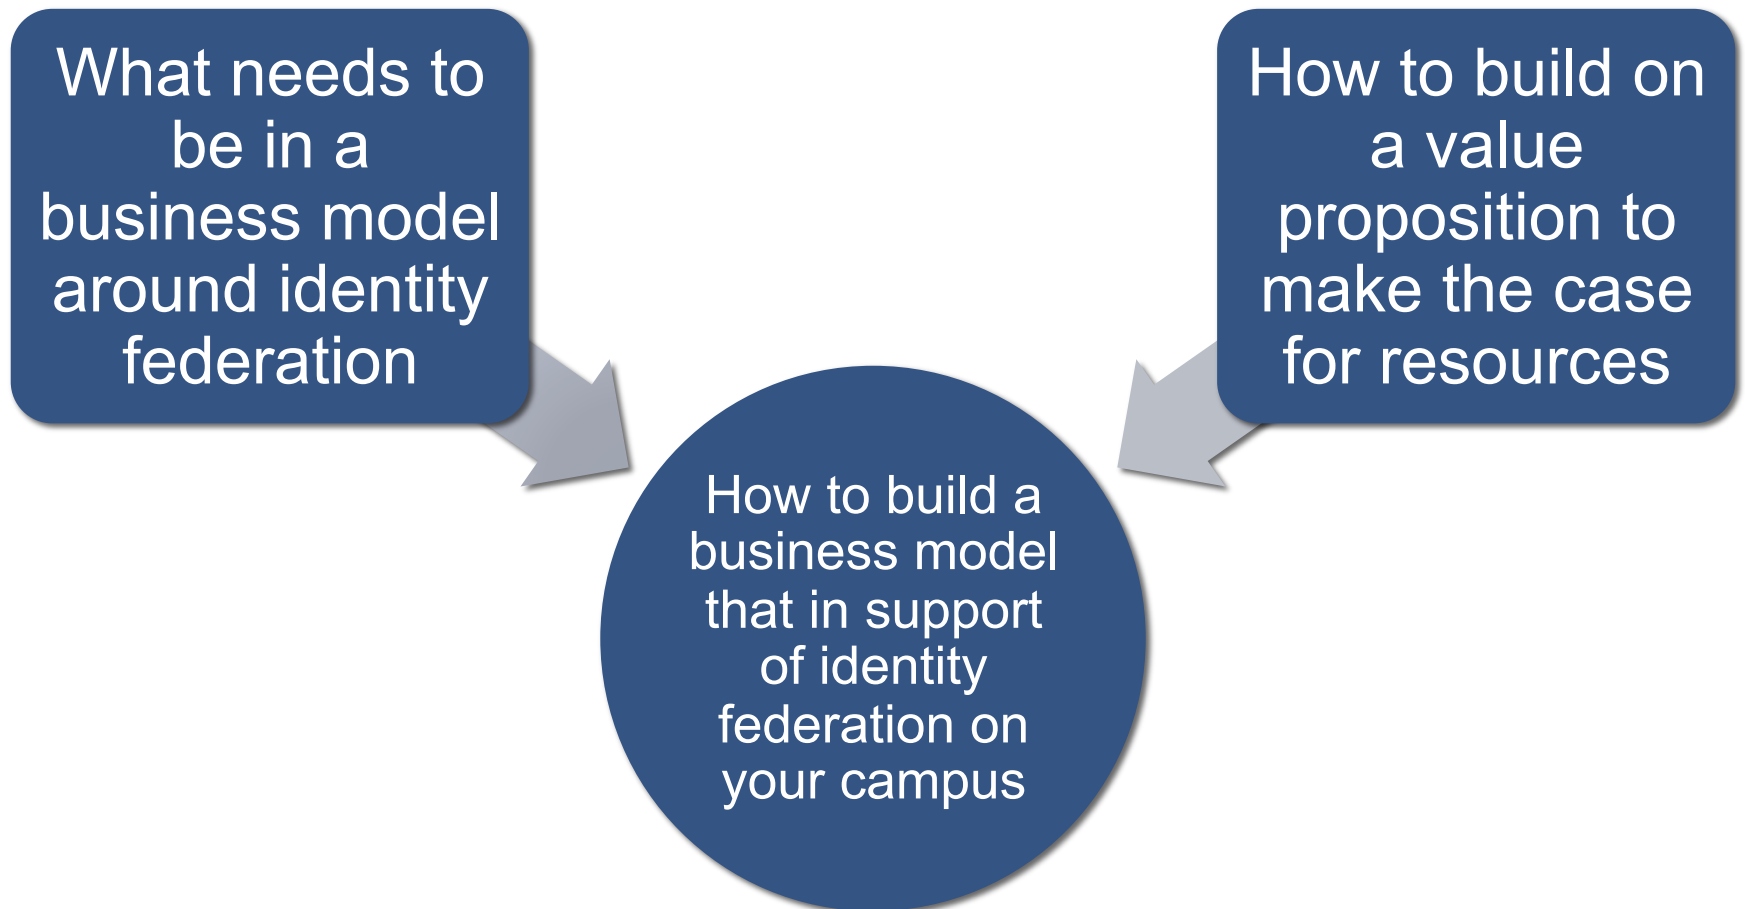
- “Ready the Pipes” – *Campus Technologies*.
<https://campustechnology.com/articles/2010/03/01/ready-the-pipes.aspx>
- “Lowering costs of identity proofing by federated identity management” – Swedish Alliance for Middleware Infrastructure.
http://www.incommon.org/docs/other/SWAMI_federated_idm_roi.pdf
- “Identity Management Toolkit” – JISC.
<https://identitymanagementinfokit.pbworks.com/w/page/50989755/Home>

BUILDING A BUSINESS MODEL

Heather Flanagan, NSRC



Learning Objectives

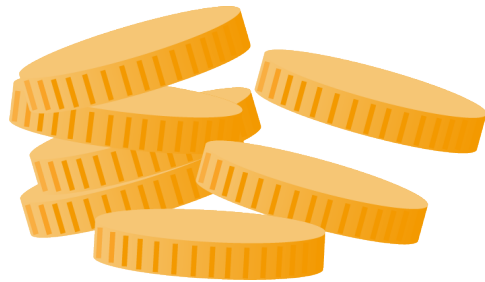


Business Context for Identity Federation

Identity is highly strategic to some commercial providers, who are trying to control the space. An anchor service such as a social network or email service often makes these providers attractive to users.

Services are also moving onto Cloud platforms that feature easy integration with the operator's AAI. This is leading to the adoption of non-interoperable AAI.

Trust is becoming a significant issue. While the network creates many positive opportunities it also introduces risks, particularly with the growth of Cloud. Users (or their organisations) do not trust some of these entities, and some actors are even considered hostile.



In an increasingly constrained budgetary environment, *fund*ers are consolidating funding on horizontal activities such as eID. Positioning and communicating our T&I work is more critical than ever; the NRENs must articulate how we add value, given these other activities.



Building a Business Model

- Learn about the technology required (in this case, the tools and platforms around identity and access management)
- Understand and develop the policies required (including federation policies, organizational policies, and security policies)
- Develop a business model regarding the operations of the federation (often the service will need to be self-sustaining; grant funds help but cannot be a long-term solution)
- Create a Service Delivery system to support the use of the service (for example, web content and a knowledge base for help desk support, training, communication and outreach, and marketing)

The Value Proposition for Identity Federations

- <https://wiki.refeds.org/x/MoA4>
- Authors: Chris Phillips (CANARIE), Lucy Lynch (NSRC), Nicole Harris (REFEDS/GÉANT), Heath Marks (AAF)
- Editor: Heather Flanagan (REFEDS/Spherical Cow Group)
- Contributors: Ann Harding (SWITCH), Klaas Wierenga (Cisco)

What is the Value?

- Collaboration Opportunities
- Reputation and Branding
- Network Security
- Budget and Business

How to Make Federated Identity Work

- Identify the business model
- Make the case
- Track metrics
- Report on the value received



“the voice that articulates the mutual needs of research and education identity federations worldwide”

<https://www.refeds.org>



The NSRC cultivates collaboration among a community of peers to build and improve a global Internet that benefits all parties. We facilitate the growth of sustainable Internet infrastructure via technical training and engineering assistance to enrich the network of networks.

Our goal is to connect people.

www.nsrc.org