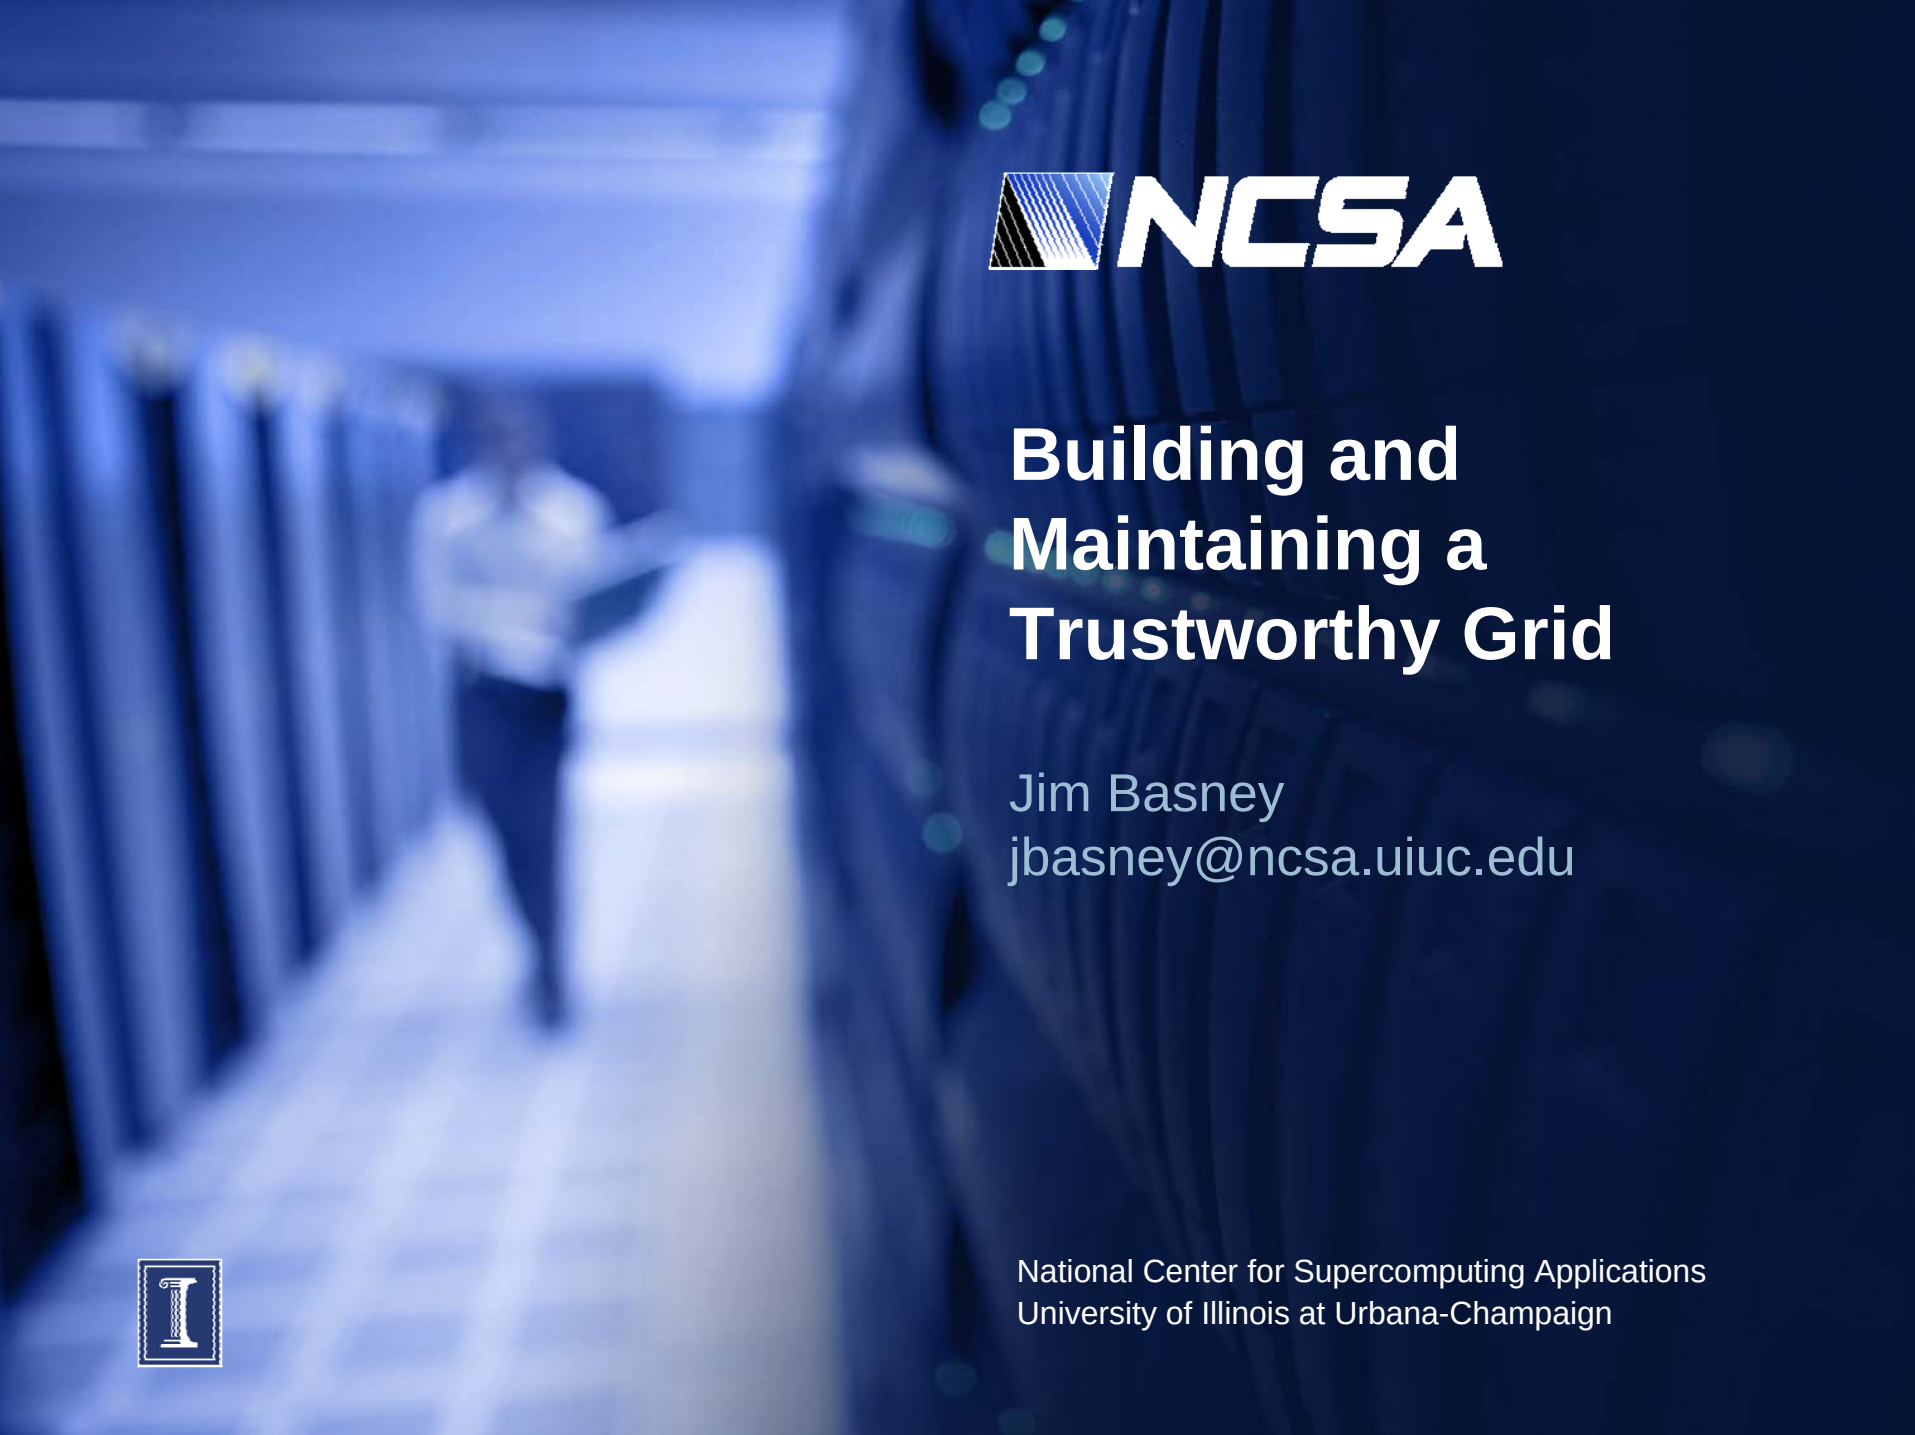

Building and Maintaining a Trustworthy Grid

Jim Basney
jbasney@ncsa.uiuc.edu

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

Outline

- Introductions
  - NCSA, TeraGrid, and Open Science Grid
- A Trustworthy Grid
- Credentialing
  - International Grid Trust Federation
- Software Vulnerability Handling
  - Globus Security Committee
- Incident Response
- Conclusions

Outline

- **Introductions**
  - NCSA, TeraGrid, and Open Science Grid
- A Trustworthy Grid
- Credentialing
  - International Grid Trust Federation
- Software Vulnerability Handling
  - Globus Security Committee
- Incident Response
- Conclusions

NCSA is…

- World leader in deploying supercomputers and providing scientists with the software and expertise needed to:
  - Fuel scientific discoveries
  - Advance the state-of-the-art in engineering
- Unique partnership among the University of Illinois, state of Illinois, and US federal government
- Home to more than 300 computing experts and students
- Key partner in the US National Science Foundation's TeraGrid
- Home to Blue Waters, expected to be the most powerful computer for open scientific research when it comes online in 2011

TeraGrid™

NCSA provides…

- More than 140 teraflops of computing power
- Consulting for 2,000 users nationwide
- Software and tools
- New insights into data through visualizations and data mining tools
- Educational programs for grade school through graduate school and beyond
- Cyberenvironments that help researchers use our systems
- Cybersecurity training and tools for cybercrime investigations

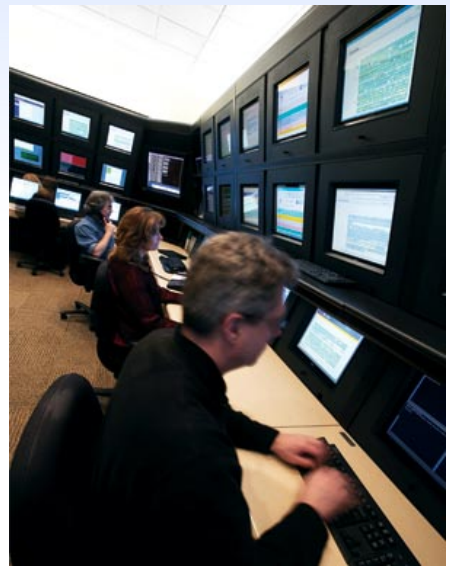

NCSA's current computing power

- 5 production systems
- More than 140 teraflops
- About 2,000 users nationwide
- Researchers receive time at no cost through peer review

Blue Waters

- Will come online in 2011
- Hundreds of times more powerful than today's supercomputers
  - Number of calculations per second sustained on real-world applications
  - Amount of memory available
  - Ability to analyze massive quantities of data
- Collaborators:
  - University of Illinois/NCSA
  - IBM
  - Great Lakes Consortium for Petascale Computation

The Blue Waters Project

- Will enable unprecedented science and engineering advances
- Supports:
  - Application development
  - System software development
  - Interaction with business and industry
  - Educational programs
- Includes Petascale Application Collaboration Teams that will help researchers:
  - Port, scale, and optimize existing applications
  - Create new applications

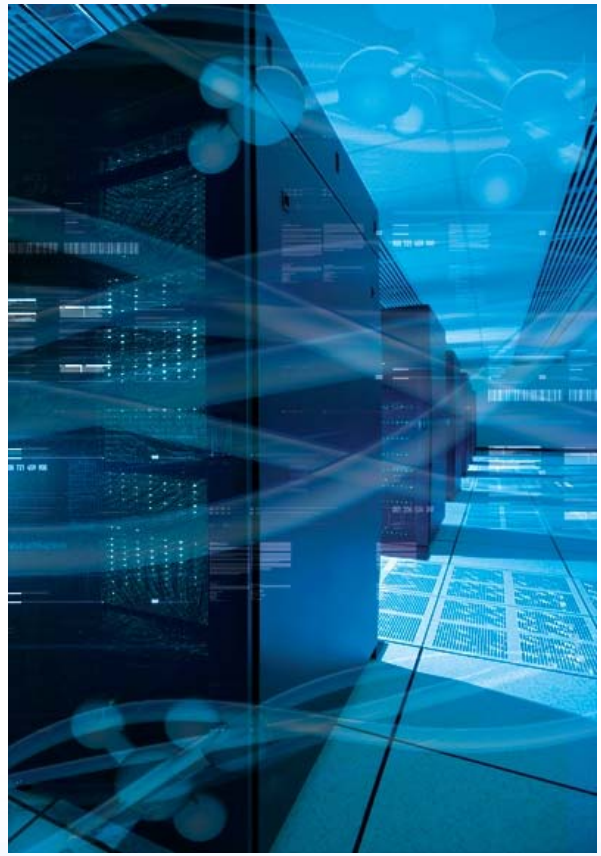

Petascale Computing Facility

- Future home of Blue Waters and other NCSA hardware
- 88,000 square feet, 20,000 square foot machine room
- Water-cooled computers are 40 percent more efficient
- Onsite cooling towers save even more energy

What is the TeraGrid?

NSF-funded facility to offer high end compute, data and visualization resources to the nation's academic researchers

(7500+ registered users from 450+ organizations)

TeraGrid™

TeraGrid Science Gateways

Use SAML assertion to convey user identifier and email address

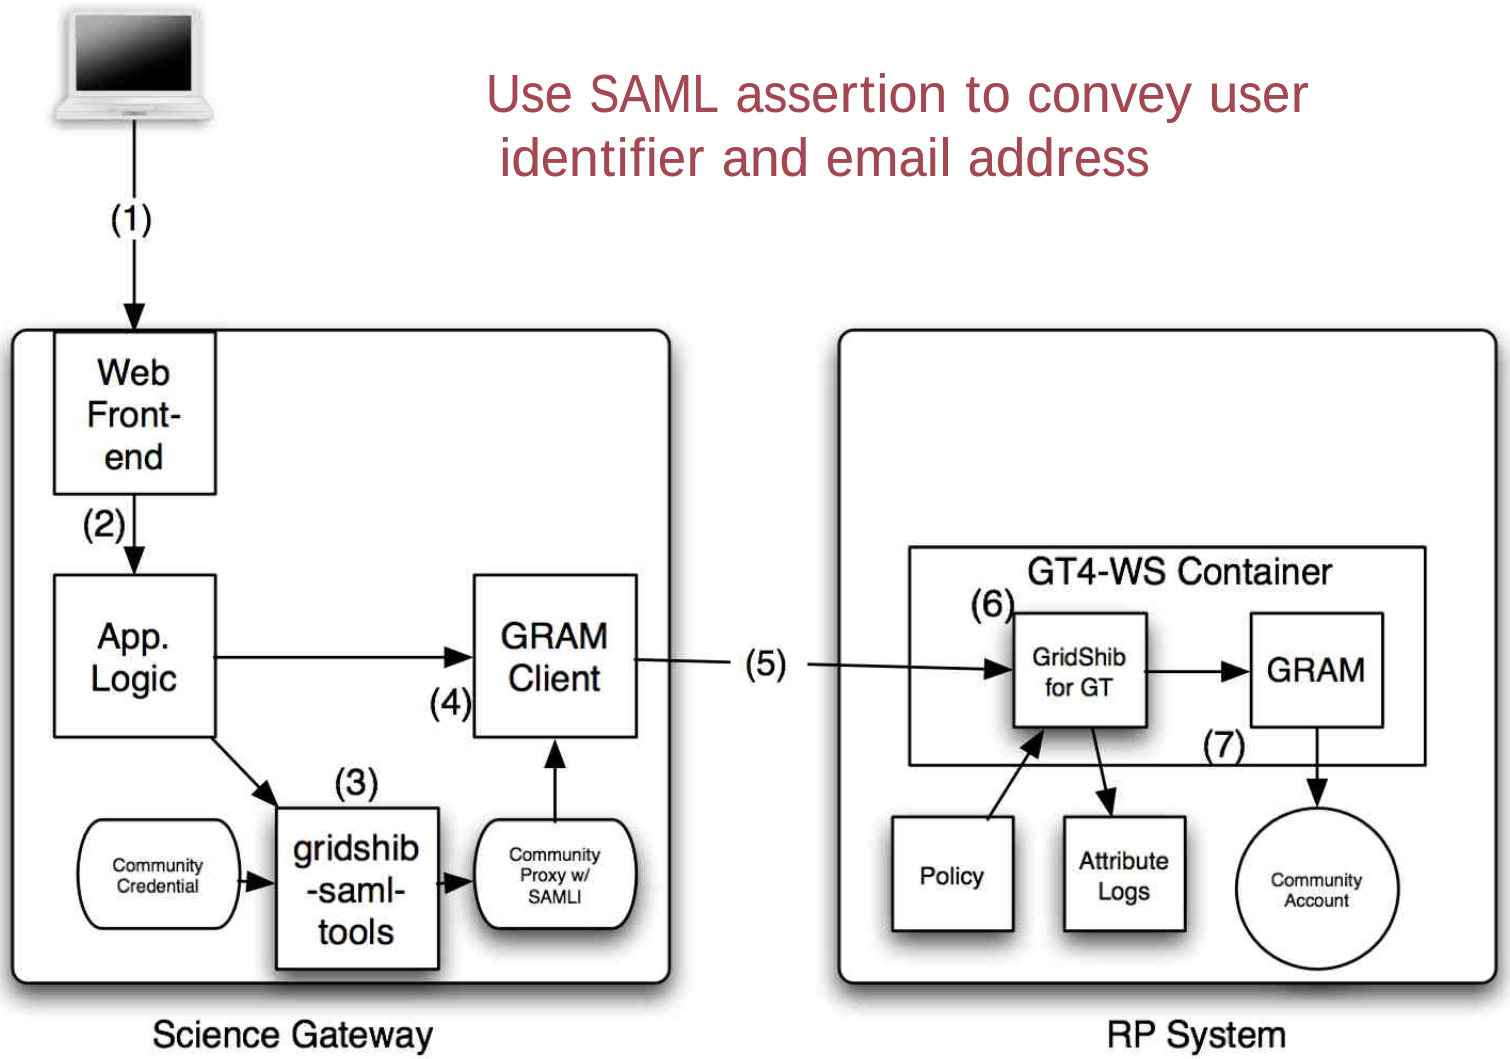

What is the Open Science Grid?

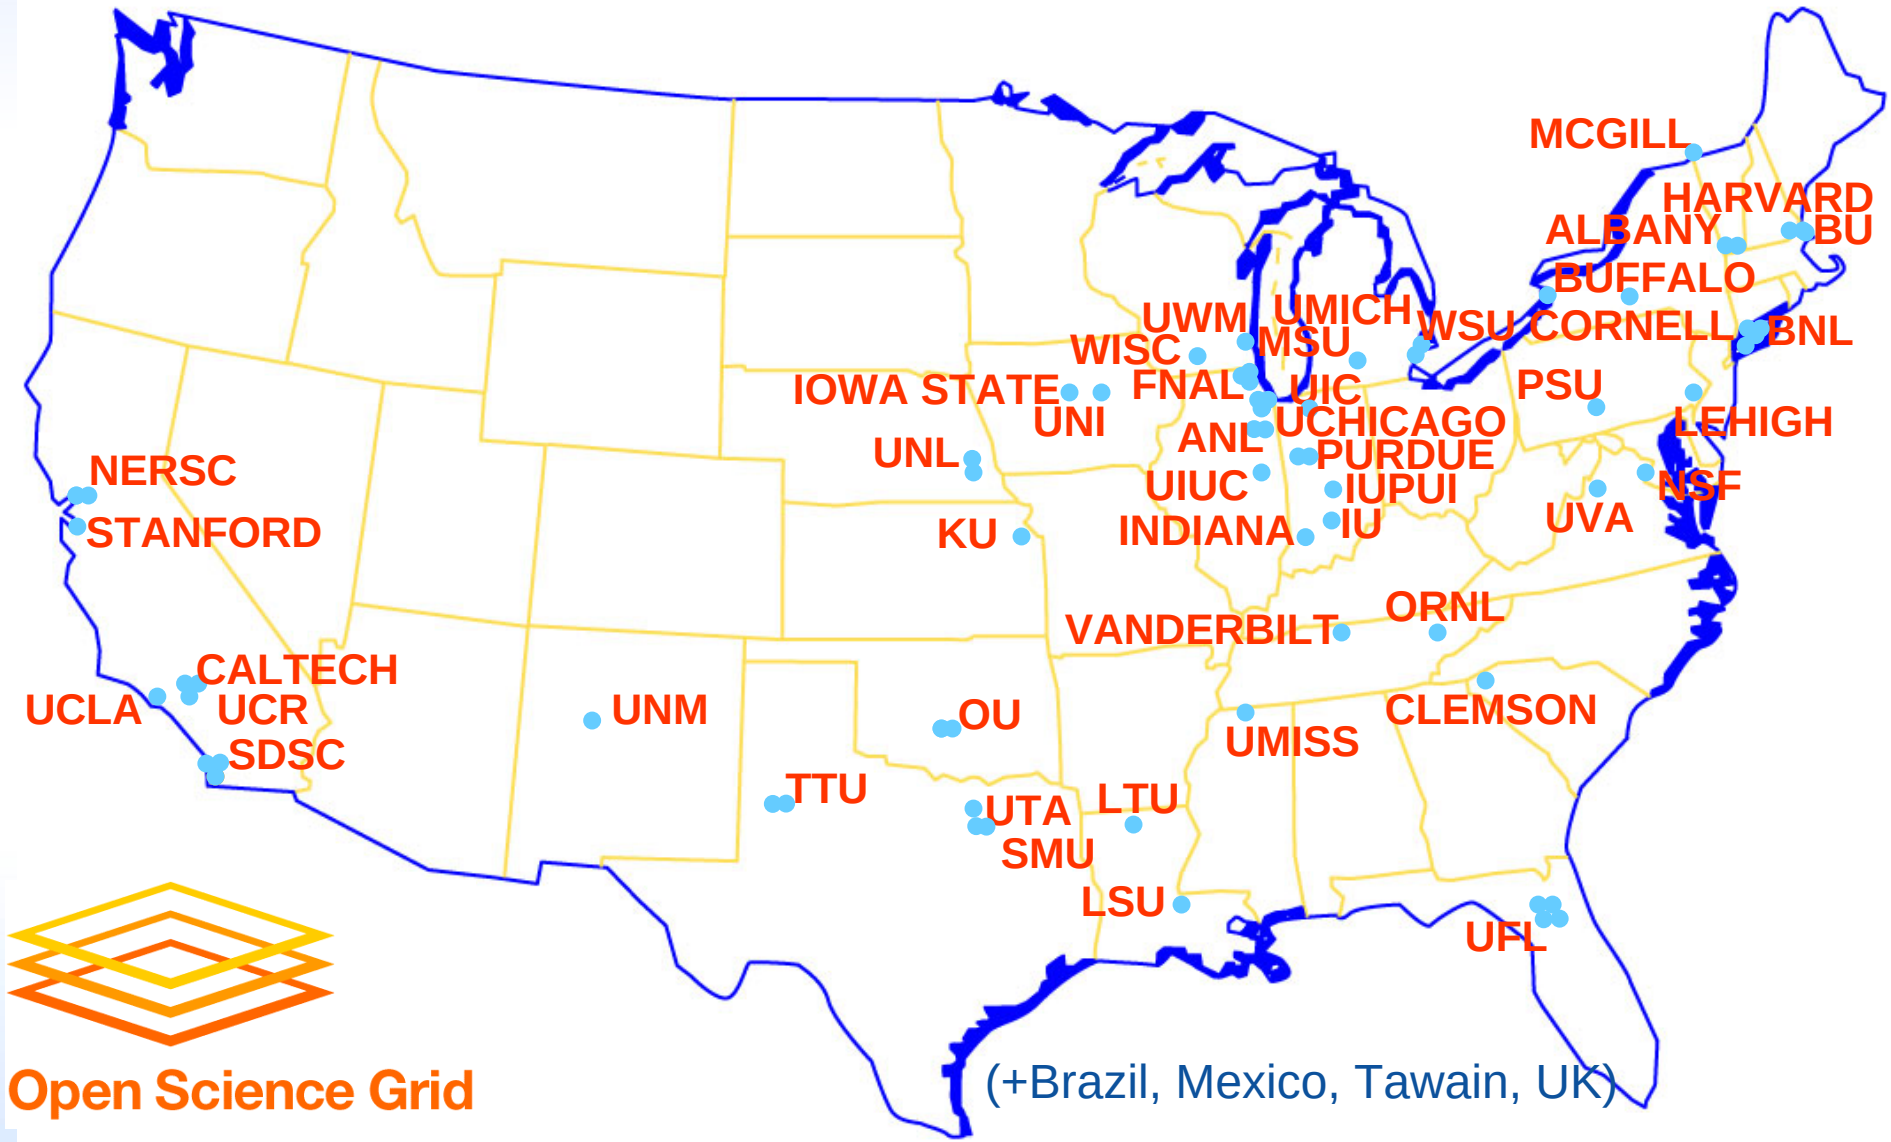

TeraGrid and Open Science Grid Compared

	TeraGrid	Open Science Grid
Number of RP sites:	11 HPC centers	5 DOE labs, 65 universities, 5 partner campus/regional grids
Resources:	24	129
Processor cores:	>100,000	>43,000
Primary workload:	Tightly-coupled parallel MPI jobs	Loosely-coupled workflows of sequential jobs
Allocation mechanism:	PIs submit proposals to national review board	Virtual Organizations (CMS, LIGO) apply to OSG for membership
Users	>7,500	>10,000
User management:	Users register directly with TeraGrid; PIs add users to projects	>30 VOs manage user membership; RPs serve VOs
Typical access:	ssh+qsub; portals	VO workload manager

Outline

- Introductions
  - NCSA, TeraGrid, and Open Science Grid
- **A Trustworthy Grid**
- Credentialing
  - International Grid Trust Federation
- Software Vulnerability Handling
  - Globus Security Committee
- Incident Response
- Conclusions

A Trustworthy Grid: Defined

- Provides reliable service to researchers
  - Avoids downtime caused by security issues
- Maintains security and privacy of research data
- Facilitates appropriate use of valuable systems
  - Acceptable Use Policy, Resource Allocation Policies, etc.
- Expands access to resources without making them less stable or more vulnerable to attack
  - New connections enabled by the grid can also spread attacks – must be addressed by security policies and procedures

A Trustworthy Grid: Challenges

- Trust is distributed and interdependent
  - CAs vouch for identities
  - VOs and Science Gateways vouch for users
  - Resources and services are interconnected
  - Incident response crosses boundaries
- Scaling to large numbers of users and resources
  - Requires delegation of responsibility, automation of procedures
- Complexity
  - Impacts usability, manageability, auditability
  - Makes problem diagnosis/response more difficult

A Trustworthy Grid: Risks

- Compromise of credentials by attackers
- Compromise of services by attackers
- Misuse of services by insiders
- Misconfiguration of services by administrators
- Hardware and software faults
- Infrastructure failures (power, network, …)
- Natural disasters
- …

A Trustworthy Grid: It's All About People

- As much as I like to talk about technology, the most important part is the people…

- Face-to-face meetings remain critical
  - Building consensus, mutual understanding, working relationships
  - Both within and between groups/organizations
- Setting standard policies and procedures
  - While allowing them to evolve
- Building support networks
  - Troubleshooting and problem diagnosis
  - Incident response coordination and information sharing
  - Sharing best practices and experiences

Outline

- Introductions
  - NCSA, TeraGrid, and Open Science Grid
- A Trustworthy Grid
- **Credentialing**
  - International Grid Trust Federation
- Software Vulnerability Handling
  - Globus Security Committee
- Incident Response
- Conclusions

Credentialing

- Issuance of credentials to users and services is a prerequisite for secure online interactions
  - Credentials prove/convey identity, attributes, memberships
- Goals:
  - Single sign-on, avoid juggling multiple credentials
  - Usability, mobility, scalability
- Options:
  - Passwords
  - Certificates
  - Assertions
  - Hardware tokens

International Grid Trust Federation (IGTF)

- A federation of regional Grid PKI Policy Management Authorities (PMAs)
  - Asia Pacific Grid Policy Management Authority (APGridPMA)
  - European Policy Management Authority for Grid Authentication in e-Science (EUGridPMA)
  - The Americas Grid Policy Management Authority (TAGPMA)
- Standards for Grid Certificate Authorities (CAs) facilitate world-wide use of grid certificates
- IGTF PMAs review and audit of Grid CA operations

www.gridpma.org

IGTF Profiles

- Identity vetting
  - Face-to-face identification of people using photo IDs
  - Verifying authorization to request host/service certificate
- Operation of CA systems
  - Protection of CA private key
  - Issuance of CRLs
  - Logging
- Credential management
  - Private keys encrypted with strong passphrases
- Interoperability
  - Certificate extensions
  - Namespace management

IGTF Risk Assessment Team (RAT)

- An IGTF subcommittee responsible for assessing risk and setting time deadlines for response and action by IGTF CAs for concerns and vulnerabilities
- Recently established in reaction to Debian OpenSSL random number generator issue (CVE-2008-0166)

https://tagpma.es.net/wiki/bin/view/IGTF-RAT

IGTF Short Lived Credential Services

- Translate a local site credential (LDAP, Kerberos) to a grid credential (certificate)
  - Leverage existing site/organization identity management
- Short-lived: certificates valid for up to 11 days
  - Local identity management updates propagate quickly
  - Integrates with local site logon
- Example: NCSA MyProxy CA
  - http://myproxy.ncsa.uiuc.edu/
  - http://ca.ncsa.uiuc.edu/

EUGridPMA Authorization Policy WG

- Addressing policy and global trust issues related to grid authorization (AuthZ).
  - Minimum requirements and best practice for the operation of a grid AuthZ attribute authority
  - Minimum requirements and best practice for VO user and service membership management
  - Accreditation of Attribute Authorities (AA)
  - Accreditation of VOs and their membership management procedures
  - Repositories and distribution of accredited AA roots of trust
  - Technical details of attribute signing and trust validation

Open Grid Forum CA Operations WG

- Develops operational procedures and guidelines for cross-grid authentication
- Produced GFD-C.125: Grid Certificate Profile
  - Referenced by IGTF profiles
  - Practical guidelines on certificate extensions, distinguished names, key lengths, etc.

www.ogf.org

Joint Security Policy Group (JSPG)

- Prepares and maintains security policies for adoption by WLCG and EGEE
  - Policies are adopted by other grids (such as OSG) for interoperability
  - Members from other grids are welcome
- Policies
  - Virtual Organization Management
  - Site Registration
  - Incident Response
  - Approval of Certificate Authorities
  - Traceability and Logging

www.jspg.org

Outline

- Introductions
  - NCSA, TeraGrid, and Open Science Grid
- A Trustworthy Grid
- Credentialing
  - International Grid Trust Federation
- **Software Vulnerability Handling**
  - Globus Security Committee
- Incident Response
- Conclusions

Software Vulnerability Handling

- The landscape:
  - Grid middleware typically combines software from many providers
  - Grid software builds on existing open source software (OpenSSL, BouncyCastle, Apache)
  - Open source development can bring developers together in an ad-hoc manner
  - Grid software often outlives original development funding
  - Individual software components are bundled into distributions (Globus Toolkit, Virtual Data Toolkit, EGEE gLite)
- The need for coordination and collaboration:
  - Writing good software advisories is hard
  - Coordination with software distributors and grid administrators avoids patch/upgrade panic

Globus Security Committee

- Community process for handling potential security vulnerabilities in Globus software
  - Provides a critical mass of responders
  - Provides a consistent, documented process for vulnerability handling
  - Participating grids have an opportunity to provide input before security advisories are publicly announced

http://dev.globus.org/wiki/SecurityCommittee

Outline

- Introductions
  - NCSA, TeraGrid, and Open Science Grid
- A Trustworthy Grid
- Credentialing
  - International Grid Trust Federation
- Software Vulnerability Handling
  - Globus Security Committee
- **Incident Response**
- Conclusions

Incident Response

- An incident is any real or suspected event that poses a real or potential threat to the integrity of services, resources, infrastructure, or identities.
- Coordination:
  - Site Computer Security Incident Response Teams (CSIRTs) are the responders on the scene
  - Inform regional/national CSIRTs, Information Sharing and Analysis Centers (ISACs), and other coordinating groups
  - Software providers consult and fix vulnerabilities
  - CAs revoke compromised certificates
  - RPs, VOs, Gateways disable accounts, apply patches, recover systems

Incident Response: TeraGrid

- TeraGrid Incident Response Team consists of CSIRT members from the 11 TeraGrid partner sites
- Single point of contact
  - help@teragrid.org
  - +1 866 907 2383
  - 24/7/365 response
- TeraGrid-wide accounts and services mean that coordinated response is essential
  - Centralized ticket tracking system
  - Emergency contact directory
  - Secure teleconference lines
  - Secure email lists

http://www.teragridforum.org/mediawiki/index.php?title=TeraGrid_Security_Playbook

Incident Response: Open Science Grid

- OSG Incident Response Team (IRT) consists of project security, operations, software, and executive staff
  - Central team coordinates with VO and site security contacts
  - Site CSIRTs not actively engaged with OSG
- Large VOs span EGEE and OSG
  - Requires coordination with EGEE IRT
  - Adoption of JSPG incident response policy
- Single point of contact
  - security@opensciencegrid.org
  - +1 317 278 9699
  - 24/7/365 response

https://twiki.grid.iu.edu/bin/view/Security/IncidentResponseProcess

SELS: A Secure Email List Service

- Used by TeraGrid incident response team
- Provides message-level security for emails exchanged on mailing lists
  - Confidentiality, Integrity, and Authentication
- Minimally trusted List Server
  - Novel feature: List Server does not get access to email plaintext
  - Proxy encryption techniques enable transformation of ciphertext
- Development with COTS and open-source components
  - Integrated with GnuPG on subscriber side; no software to install
  - Integrated with Mailman on server side with easy installation and setup

http://sels.ncsa.uiuc.edu/

Conclusions

- Trustworthy grids require a strong security community
  - Groups: IGTF, JSPG, Globus Security Committee, OGF CAOPS
  - Building on community standards and best practices
- Trustworthy grids provide reliable service
  - Security should enhance (not obstruct) quality of service
- Different types of grids require different security practices
  - Comparison between TeraGrid and Open Science Grid
- Trustworthy grids rely on worldwide collaborations
  - Software providers, resource providers, incident responders, credential providers, …, and most important: researchers who use the grid

Thanks!

- Stuart Broughton / University of Canterbury
- Grid security experts around the world
- TeraGrid and OSG staff

This material is based upon work supported by the United States National Science Foundation and Department of Energy. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation or Department of Energy.