Grid Security

Dane Skow

Deputy Computer Security Executive, Fermilab

SEC Area Director, GGF

Masterclass Gridforum.nl

May 18, 2004

Outline

- What I will talk about
  - GGF organization and goals
  - Grid Security Architecture
  - Current hot topics in Grid Security
- What I won't talk about (except Q&A)
  - Crackers and their goals
  - Relative merits of various security tools
  - THE SOLUTION

Fermilab

- Fermilab is a US Department of Energy (DOE) National Laboratory devoted to high-energy physics and astrophysics.
  - Highest Energy Collider in the world
  - ~12,000 computers of many sorts
- We host current experiments with collaborating institutions from all over the globe (including NIKHEF)
- We are a Tier 1 facility for the LHC focussing on the CMS collaboration.
- A Founding member of the Open Science Grid effort forming now in the US.

Structure and mission of GGF

- Global Grid Forum
  - Forum for development of Grid communities
    - Develop Grid vision
    - Advance distribution of technology
  - Standards organization
    - Advance interoperable technologies
    - IETF has IP as the common base
    - W3C has XML and HTML
    - OASIS builds on XML from the bottom up.
    - GGF aggregates Grids from the top down

GGF Structure

GGF Corp.

Secretariat: Operations

- Board of Directors
  - (Catlett, Messina)
  - •Financial, Legal Responsibility
  - •Appoints GGF Chair
- Conference Mgmt
- Sponsor Programs
- Finance/Legal
- IT (Website, etc.)
- Staffing & Services

Holds *non-exclusive* copyright for document series

GGF

Document and Standards Work

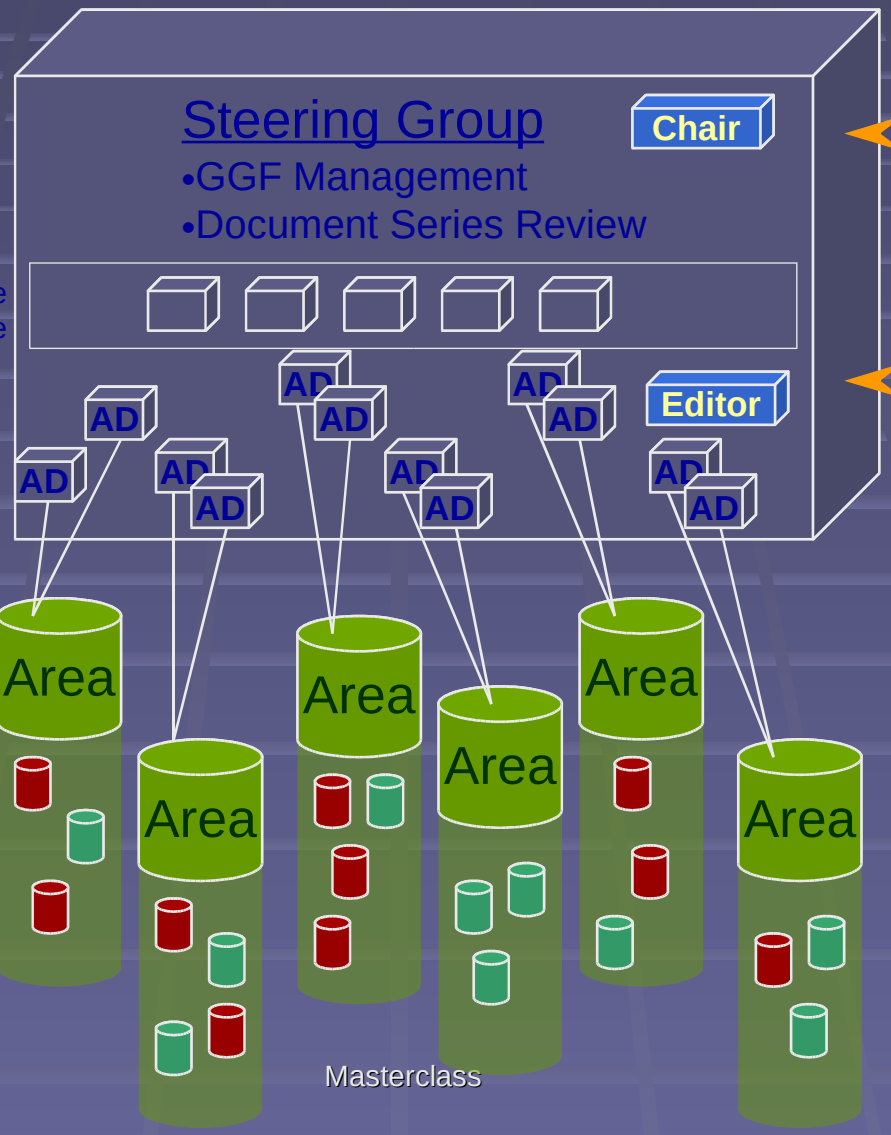

- GGF Advisory Committee
  - •Advise on strategic direction, industry, government
  - •Advise on governance
- Grid Research Oversight Council
  - •RG/WG Workshop series
  - •Advise on Research Issues

= working group

= research group

GGF Steering Group (GFSG)

- GGF Steering Group
  - Charlie Catlett (ANL) [Chair]
  - Ian Baird (Platform Computing)
  - John Tollefsrud (Sun)
  - Peter Clarke (UCL/UK)
  - David Martin (IBM)
  - Cees DeLaat (UVA/NL)
  - Andrew Grimshaw (UVa/Avaki)
  - Marty Humphrey (UVa)
  - Dane Skow (FNAL)
  - Craig Lee (AC)
  - David Snelling (Fujitsu)
  - Bill Nitzberg (Veridian)
  - Jennifer Schopf (ANL)
  - Satoshi Matsuoka (Tokyo Inst. Tech)
- "At-Large" GFSG Subcommittee
  - Ian Foster (ANL/UC)
  - Bill Johnston (LBL)
  - Ken Klingenstein (Internet2)
  - Dennis Gannon (IU)
  - Alan Blatecky (SDSC)
  - Jeff Nick (IBM)

- GFSG Role
  - Operational Management and policy
  - Document series review
  - Chartering of new groups
  - Group oversight and review
- GFSG Structure
  - Two area directors per area
  - At-Large Subcommittee
    - Appeals Process
    - Oversee liaisons with Other Groups
      - IETF, Internet2, W3C, DMTF, OASIS, IPv6 Forum
- GFSG Participation
  - Individuals, not representatives
  - Selected by Nomcom

Sec focus areas in GGF

- Fundamental Design Frameworks
  - Framework documents
    - Site AAA (completing)
    - Authorization Frameworks (completing)
    - OGSA Security
- Operational Experience of existing Grids
  - CA operation
- Standardization
  - OGSA authorization

Glossary

- Identity: a unique way of identifying an actor on the grid. Implies a namespace control system.
- Authentication: determining the identity of the actor making the request
- Authorization: determining if the request is permitted.
- Auditing/Accounting: being able to associate actions with requestors in a reliable fashion
- Last 3 frequently referred to as AAA

Architecture vrs. Blueprints

- Architecture is definition of the essential elements which define a style.
  - Are Grids for science and business the same architecture, style, or even consistent ?
- A blueprint is a design in enough detail to allow independent builders to interact to create a coherent implementation.
  - These are the specifications which are of particular value when mass producing and/or coordinating multiple “subcontractors”

What is “security” ?

- A feeling of assurance ?
  - Rather like an insurance policy you hope never to use, but have (probably overblown) expectations of help in times of trouble.
- Preventing bad things from happening ?
  - Rather like a vault in which to store treasure
- A plan for what to do when bad things DO happen ?
  - Rather like the Red Cross emergency response plans
- The ability to enforce particular policies ?
  - Rather like a police capability to break up mobs

Security Requirements

- There have been several efforts in last couple years to extract the security requirements of various communities.
  - No definitive list possible.
  - Inherent need for compromises.
- GFD-12 & 18
  http://www.ggf.org/documents/final.htm
- GGF SiteAAA Research Group - https://forge.gridf
- WS-I Basic Security Profile Scenarios
  http://www.ggf.org/documents/final.htm

- "If a bad guy can persuade you to run his program on your computer, it's not your computer anymore."
  - – Microsoft Security Law #1

Architectural Elements for Security

- Sources of Identity
- Association Method
- Secure Communications
- Control Points and Responsibilities
- Secure Logging
- Distributed Authorization
- Ability to suspend operations and/or disassociate
- Contracts and/or “court of appeals”

Associations

- Users collaborate on several scales
  - Individual associations (2 users)
    - Emphasis on speed and ease
    - Want to leverage existing infrastructure
    - Collaboration tied to individuals
  - LHC Experiments (2000 users)
    - Many peers
    - Indirect support of resources
    - Collaboration must survive membership changes

More Associations

- SETI (millions of users)
  - Distributed Computing:
    - one source of application,
    - contribution of fungible resources.
- “Virtual Computers”
  - Partnership arrangement of service providers behind a common customer interface

Secure Communications

- Secure against data insertion
  - No command insertion
  - No session hijacking
- Secure against disruption
  - Withstand DOS attacks
- Secure against rogue users
  - Disaster recovery
  - Authorization control

Secure Communications

- Secure against application exploit
  - Ability to detect compromised applications
- Secure against compromise of secrets
  - Ability to restore good state of secrecy
    - Speed
    - Breadth of compromise

Control points for Security

- Management of resources
  - Collaboration management needs way to allocate resources among participants
  - Resource managers need methods of suspending resource to allow maintenance activities.
- Containment of damage
  - Principle of least privilege
  - Throttles applied for “runaway” activity

Control Points II

- Access Control
  - Virtual Organizations control membership and roles within organization.
  - Resource Providers impose access control requirements on resources, even visibility of resources.
  - Users have access control requirements on data.

Control Points III

- Troubleshooting/Forensics
  - Need a consistent audit trail
    - For system debugging
    - For application debugging
    - For incident forensics investigations
- Startup/Cleanup/Recovery
  - Frequent need to stage recovery of elements
  - Need way to clean system from aborted/failed jobs

Moving beyond GT2

- Globus Toolkit 2 is by far the most commonly deployed software base for Grids today.
- Currently de facto standard expressed as open software rather than standard specifications.
  - Hacked openssl as foundation of GSI.
  - EDG, VT, and GT gatekeeper authorization callouts
- Need to allow for ways for developer pool to expand and for “profit” to be made.

Standards for interoperability

- Necessary for connection to “legacy” resources
  - Need to allow for “adapter-ware” so that operational facilities can be brought to the table
- Necessary for collaborative partnerships
  - We need to develop a model for demarcation. Grid/Web Services are contender
- Necessary for competitive development
  - Existing specs reduce the entry costs
  - Developing specs tests the stamina

What Next ?

- Depends on who's paying.
- Industry
  - Much industrial money on Web Services
    - Commerce over current interactive Web is BIG business (estimate ?)
    - Next effort is to make an ecommerce infrastructure people can build on like the Internet.
    - Competitive markets, associated sales, etc.

Government Funding

- Government programs focus on big science projects.
  - Desire for big win like Web
  - Best chance for “jobs projects”
- Politically need to show some local advantage
  - Tendency to “embrace and extend”
  - Concerns about global collaboration
- Valuable to show commercialization
  - Best to show relevance to gov’t programs

Power to the People

- Open Source
  - Literature continues to grow in scope and quality
  - Developer community continues to grow
  - Clear, simple, open standards
- “peer to peer”
  - Great demand for global sharing of files
  - Often in conflict with IP holders’ interests
  - Global catalogs and efficient network utilization
- political association
  - Growing desire to supplement (confirm) official news sources
  - Community formation and privacy

What is the foundation ?

- IPSec functions at the IP layer
- TLS/SSL functions at the Transport Layer
- GSI Builds on top of that at the Application layer
  - Adds credential delegation
  - Needs to add richer authorization structure
- More thorough discussion at http://home.fnal.gov/~dane/RJC_DDS_2003

One-way Functions

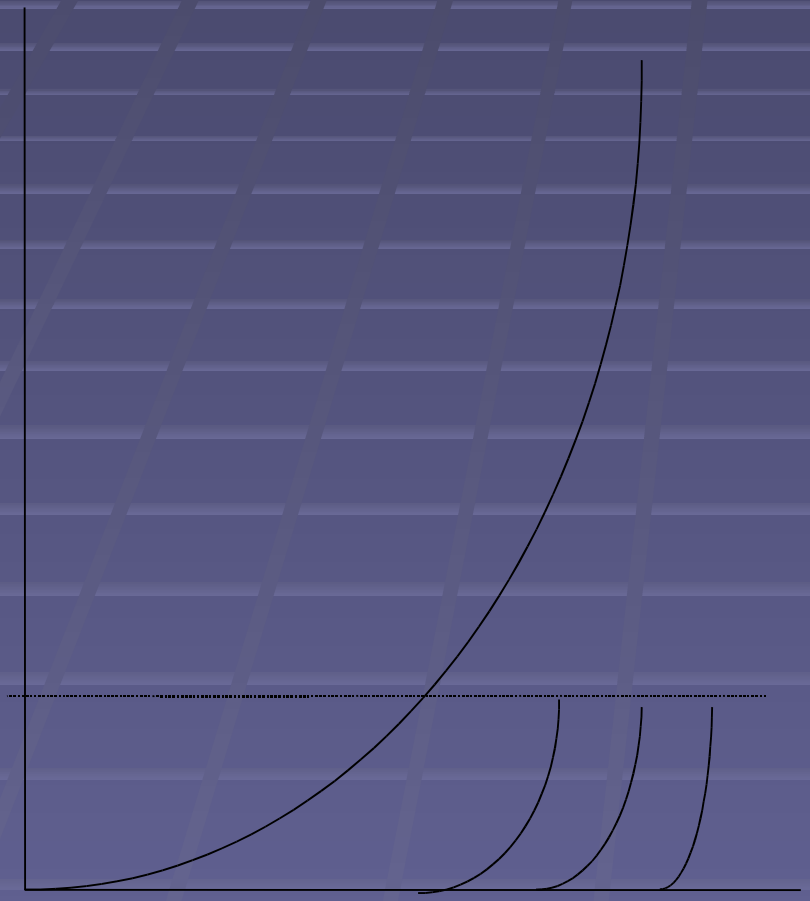

- Consider Y^x mod P as a function
- Computationally very expensive to reverse
- Alice and Bob agree (publicly), Y=5, P=7
- 5^x mod 7

Diffie-Hellman Key Exchange

- Alice picks secret number, say x=2
- $Y^x \bmod P = 5^2 \bmod 7$
- 5*5 mod 7 = 4
- Send a=4 to Bob
- b*b mod 7 = 1
- Shared secret! -- can be used as a key

- Bob picks secret number, say x=3
- $Y^x \bmod P = 5^3 \bmod 7$
- 5*5*5 mod 7 = 6
- Send b=6 to Alice
- a*a*a mod 7 = 1
- Shared secret! -- can be used as a key

Diffie-Hellman Key Exchange

- Eve intercepts values used for Y and P but can't use them to deduce x in a simple way.
- Developed and first publicly demonstrated in 1976
- Alice and Bob no longer have to meet or trust a 3rd party for key exchange
- Still inconvenient -- "real-time" exchanges to establish a key

Hot Issues in Security for Distributed Computing

- Management of Secrets
- Error Handling (Incident Response)
- Distributed Authorization
- Usability
- Identification and Privacy

Management of Secrets

- Authentication relies on one of two things: a secret, or a secure token.
  - Secrets can be exposed
    - It is notoriously difficult to know when that's happened.
    - Best practices need to be developed to minimize the likelihood of secrets being exposed in real-life usage scenarios. User education and acceptance essential.
    - Methods for rapidly and securely replacing the secrets and thus restore a good state are required.

Management of Secrets (cont'd)

- Tokens can be forged or duplicated
  - Biometrics is plagues by this problem. How do you replace a duplicated fingerprint ?
  - Information is not a tightly controlled secret. The method of presenting the information is assumed to be difficult to forge (Fingerprints)
  - Tsutomu Matsumoto able to fool many fingerpringt scanners with simple casts. http://www.schneier.com/crypto-gram-0205.html

Error Handling

- “A cluster is an excellent error amplifier.” C. Boeheim, SLAC
- “A grid is an automated error amplifier.” corollary by D. Skow

- The most likely source of early widespread denial of service is accidental misuse from an authorized user.
- What controls are in place for rooting out resubmitting jobs ?
- LCG has draft incident response plan

  http://proj-lcg-security.web.cern.ch/proj-lcg-security/documents.html

What might a bad guy do ?

- Copy credentials to go poach resources.
- Steal resources for unauthorized uses
- Disrupt operation for fun or profit
- Hunt for information
- Impersonate you to use your good name
- Modify data for fun or profit

Incident Response

- Incident Scenario #1
  - Grid worm takes initial credentials and tries to access Grid resources. On success “phone home”, scan the accessed machines for proxies, certificates and private keys, any found are used to seed the next worm instance.
- Issues:
  - Load due to Grid scanning
  - Server DOS due to session startup overheads
  - Many compromised credentials
    - How does replacement process work (user driven) ?
    - How do you associate compromised credential(s) with compromised host(s) ?
  - What to do about compromised proxies ?

Incident Scenario #2

- Grid software vulnerability is found in standard grid protocols. Worms spread across the internet attacking all available resource providers.
- Issues:
  - Network load due to worm attacks
  - Patched Server DOS due to session startup overhead ?
  - What network restrictions are needed for Grid Services ?

Distributed Authorization

Usability

- Mobility
  - The primary interface is moving to the laptop and researchers expect to use it anywhere and everywhere all the time.
  - A large business opportunity is using Grid technologies to make network ubiquitous.
- Skills
  - Need to improve training for new users to become functional
  - Xgrid an important recent entry (uses no web services, GT, etc.)

Usability (cont'd)

- Intent Description
  - Need to be able to simply describe the job that needs to be run such that it's rigid enough to be automated yet simple enough for non-experts.
  - Need to be able to simply restrict delegation authorities to retain control yet stay simple enough to be done (resist * syndrome)

“Who's responsible for this mess ?”

- Nothing is as galling as interrupting your plans to fix somebody else's problem.
- Current maintenance of CRLs, certificate expiry, gridmapfiles, etc. has a high confusion factor and maintenance load on users, sysadmins, and infrastructure support.
- Concerns that incident response load will increase
- Need to work through agreements to partition the responsibilities and develop working patterns between Grid elements.

Privacy vrs. Federated Identities

- (Need pointer to Shibboleth)
- Resource managers have increased concerns about knowing who's using their facilities.
  - Distributed process requires distributed trust.
  - Desire strong traceable attachment of an electronic identity to a person.
- Users have reasons to keep multiple identities and to keep their electronic lives private and well protected
  - Growing frequency and damage of identity theft

General discussion