Unique Challenges for Incident Response in a Grid Environment

James J. Barlow

<jbarlow@ncsa.uiuc.edu>

Head of Security Operations and Incident Response

Aashish Sharma

<aashish@ncsa.uiuc.edu>

Security Engineer

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

Overview

- Grid computing
- TeraGrid
  - Security working group
  - Security concerns
    - Vulnerabilities
    - Software development
  - Grid Incidents
  - Some TeraGrid IR Solutions

Incident Response Overview

- Goal is to minimize the impact of an incident to an organization
- Incident response steps
  - Preparation
  - Identification
  - Containment
  - Eradication
  - Recovery
  - Follow-up

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

What is grid computing?

Grid is a type of parallel and distributed system that enables the sharing, selection, and aggregation of geographically distributed "autonomous" resources dynamically at runtime depending on their availability, capability, performance, cost, and users' quality-of-service requirements.

Overview of the TeraGrid

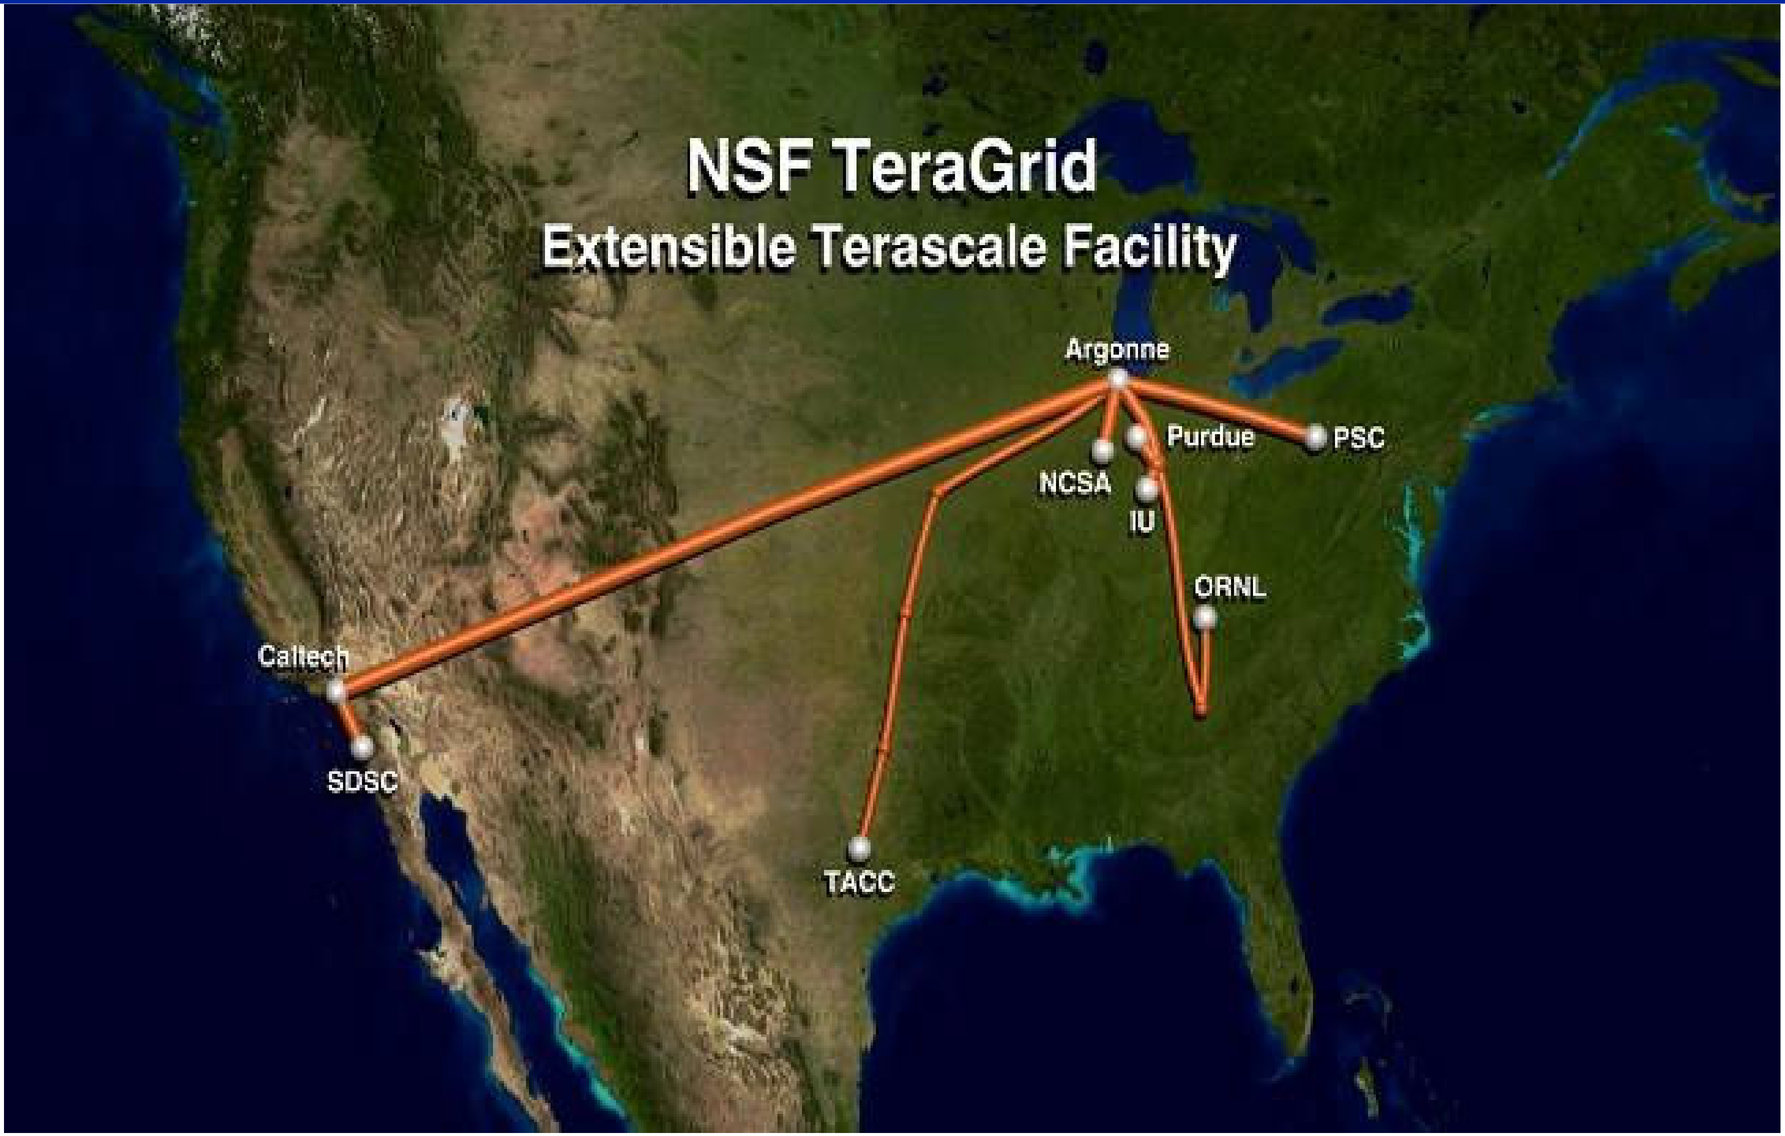

TeraGrid Mission

To provide integrated, persistent, and pioneering computational resources that will significantly improve our nation’s ability and capacity to gain new insights into our most challenging research questions and societal problems.

What makes grid computing different?

- Same userbase across multiple sites
- Global userbase
  - We don’t control the endpoints
- High profile targets
  - Zero in on the management, login, storage nodes
    - Management node is a big target
  - Monitoring node
- Homogenous systems
  - Hardware
  - Software
    - CTSS (Common TeraGrid Software Stack)

Other security considerations

- All machines on public address space
  - Why? (grid ftp, cluster nodes access other clusters)
- Administrative issues
  - Separation of privileges
  - No direct root logins
    - OTP for root escalation
- Each site does not do the same level of monitoring
  - Flows, IDS, syslog correlation, File integrity checks

TeraGrid Security Working Group

- Policies
  - TeraGrid Newbie guide
  - Security Memorandum of Understanding (MOU)
  - Security Playbook
- Bi-weekly security-wg calls
- Weekly incident calls
- Email lists
  - Security-wg list
  - Incident lists
  - Encrypted emails
    - Looking into SELS (http://sels.ncsa.uiuc.edu/)

What kind of vulnerabilities have we seen?

- Globus
  - Open source software toolkit used for building grids
- Globus Security Advisories 2006-1 and 2006-2
  - How temporary files are handled impacted MyProxy service
- Globus Security Advisory 2007-02
  - GSI-OpenSSH vulnerability

the globus toolkit

Vulnerabilities (cont.)

- Gx-map
  - Automates the maintenance of files in the /etc/grid-security directory
- Earlier versions found to allow a malicious user to insert arbitrary grid-mapfile entries
- Why do we worry about these?
- Scope is beyond own organization

Other “software developer” worries

- HPN-SSH (High Performance SSH/SCP)
  - Developed at Pittsburgh Supercomputing Center
  - Allows cleartext transmission after authentication is negotiated
  - Saw up to 80x speed increase in transmissions
- Patch now in all TG systems
- Problem: Developers (and some users) wanted cleartext as the default setting

Developer worries (cont.)

- Portals (or science gateways)
  - Central place for a community to work and submit jobs
  - Can act similar to a group account on a machine (called community accounts)
- Accounts created on all TG resources
- Security concerns
  - Who developed the portal
  - Who manages the portal
  - How is the portal handling security issues?
    - Audit trail: Who used the resource and when?
    - Credential management
  - Each RP may have individual concerns

Grid incidents

- User account compromises
  - Most common problem so far
  - Don’t control the end user’s systems
- Seen used for a number of malicious things
  - Scanners
  - SPAM
  - DoS
  - Joins IRC
- Can sometimes be much worse if there are vulnerable systems within organization

Grid incidents (cont.)

- **Network monitoring node**
  - **Each TG site has one**
  - **Network admin desktop compromised**
    - **Allowed access to monitoring node**
    - **Critical network location**
- **Started running Cain & Abel**
  - **Sniffed SSHv1 information, other cleartext passwords**
- **Accessed TG network working group's secure server**
- **Tried social engineering admins over IM**

Social engineering attempt

[19:18] NM: can you arrange me short term root on tglogin? We've had a breach.

[22:04] pld: which tglogin and why?
[22:04] NM: huh?
[22:05] pld: [19:18] NM: can you arrange me short term root on tglogin? We've had a breach.
[22:05] NM: hahha
[22:05] NM: that wasnt me
[22:06] NM: well here we go

Social engineering attempt (cont.)

NM: how's it going?
PT: prety good... how about you
PT: killed any kitties recently
NM: fubar as usual
NM: lol
NM: not today
PT: <social comment>
NM: ack
NM: ok i guess
NM: has the root password been changed to tglogin? There are some rogue processes (allegedly) and I can't get access
NM: you still there buddy?
PT: Here... let me just send you the root PW
PT: :)
PT: I am not sure if we are still buddys.. we never talk
NM: don't be coy

[chat snipped]

NM: did you send that cs to my email? if not can you send it to my yahoo account.
PT: cs?
NM: i meant tg root
PT: I was just going to IM is
PT: is=it
NM: yeah that's probably safer
PT: I'll just post it on my blog
NM: ...
NM: sorry to be an ass but can you help me out or not?
PT: with what?
NM: tg?
PT: you really need th root password for some tg node?
PT: No I can not help you... that isn't my machine
PT: I assumed you were joking
NM: i just need to look at tglogin for like 10 seconds, you can do it yourself if you don't trust me
PT: tg-login1 at NCSA?
NM: preferably

Social engineering attempt (cont.)

PT: what do you want me to lok at?
NM: we think there's a rootkit running
PT: wow
NM: i tried to sound urgent earlier
PT: sorry
PT: what should I look for?
PT: (I am not socialy capable of giving you root)
NM: an sshd with a ps starting with 4444 - I THINK that's the sig

[chat snipped]

PT: what makes you think it is comprimised
NM: i've seen the same thing here on d*
PT: what?
NM: well
NM: we got owned pretty badly. - that's classified btw
PT: O

Real user now back online sometime later:

PT: I need to get going...
NM: yes
PT: If you think there is a chance that we are in trouble, you should alert someone here
NM: who are you talking to?
PT: you
NM: trouble?
PT: with root
NM: i have no idea
NM: i will if i know anything
PT: ok
PT: thanks

Large Scale Grid Incident

- Large scale attack across the TeraGrid
  - Some may have heard of it as the “TeraGrid incident”
- Targeted login nodes to trojan ssh client
  - Had a number of tools to try compromising sites
- Very persistent
  - Not very many automated tools
- There are a few reasons attributed to his success
  - Sites generally trusted their users
    - Generally lax at patching local exploits
  - Effectively utilized accounts across multiple sites
    - Always checked known_hosts files
  - First time we had seen a grid related attack

Large Scale Grid Incident (cont.)

TeraGrid IR Solutions

- **Emergency contact lists**
- **Incident response forms**
  - **User account questionnare**

```
- Do you use the password of the account at other TG sites or other general accounts
  (Hotmail, Amazon, Paypal, Ebay)?
- What was the time of your last known login?  Where was it from?
- From what locations do you usually login (hostnames/IP)?
- Which sites/machines have you used?
- Which do you expect to use?
- What locations (hosts) can we expect to you to login from?
- Can accounts at other TG sites be closed down, or do you expect to use them in the future?
  If so, which sites are not needed:  (PSC, SDSC, NCSA, ANL, Purdue, Indiana, ORNL, Texas, etc.)
- Are passwords needed on all the sites, or are you using grid auth or ssh keys?
- Since the account was compromised, are there any special concerns on the data there?
  private data? grid certs?
- Do you have any idea how someone may have gotten your login info (login/passwd)?
  what machines may possibly be compromised?  your desktop? some other machine you used?
- Have you heard anything from any of these sites on hacker activities?
- If possible, please provide contact information for you local Security operation.
```

  - **Want to make sure host gets cleaned**

TeraGrid IR Solutions

- **Problem: Different usernames at different sites**
  - **Solution: TeraGrid user account lookup page**

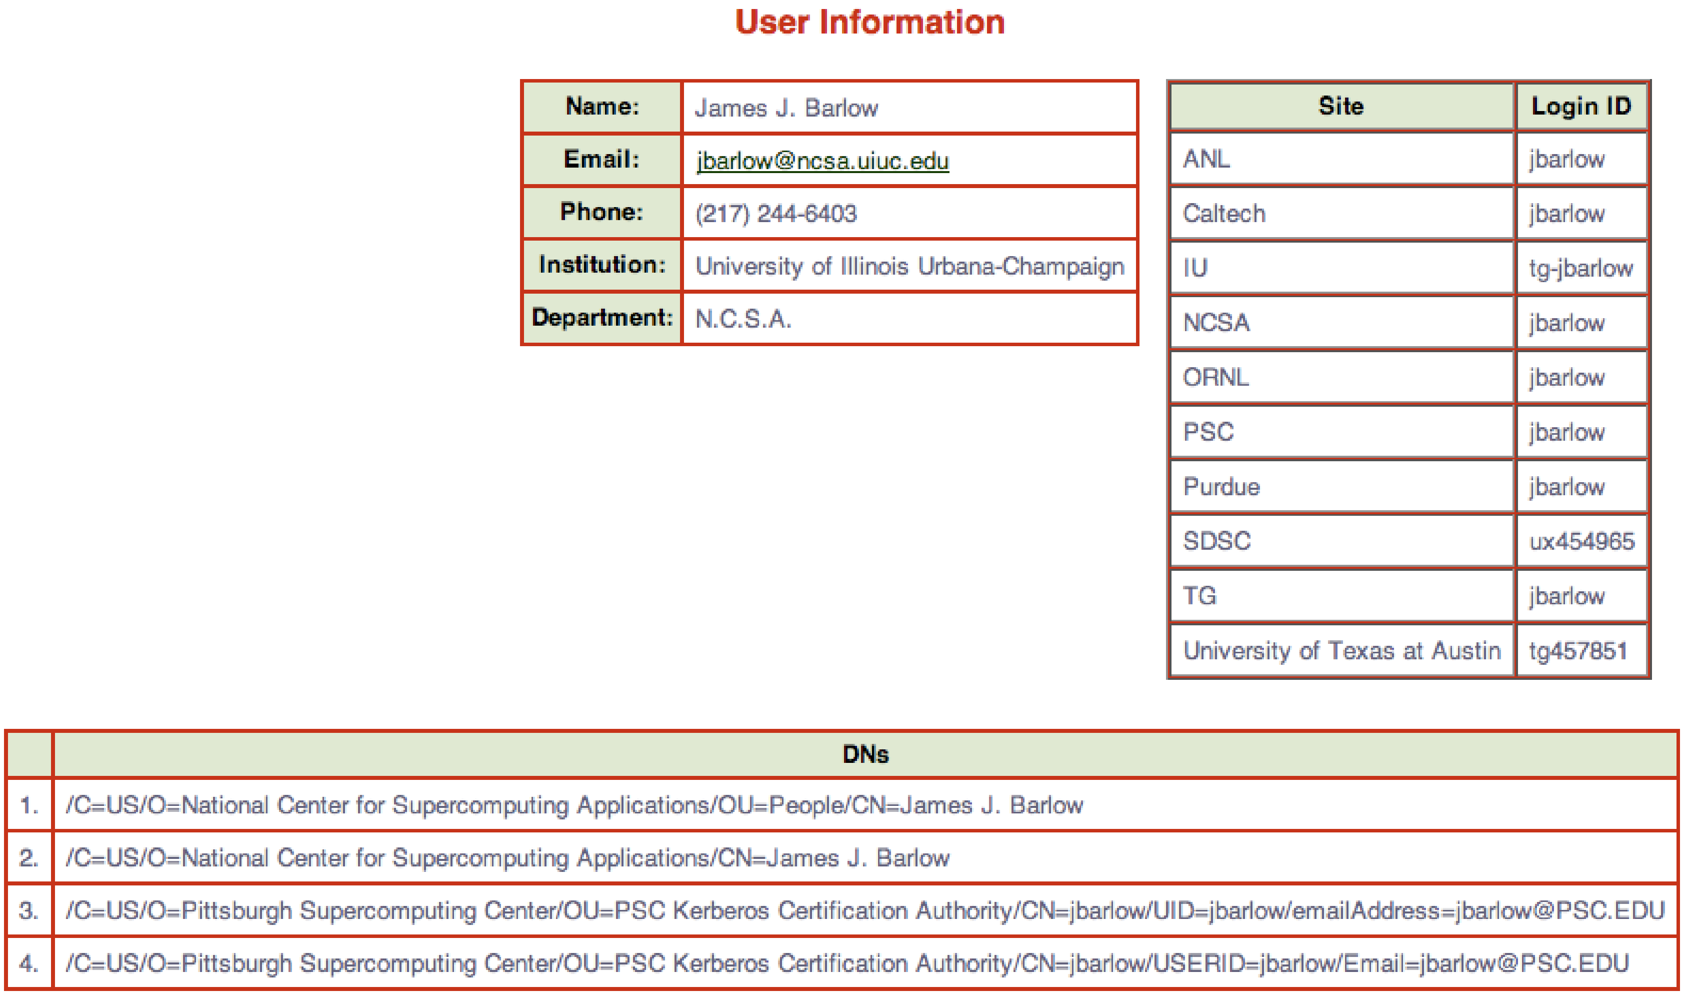

User Information

Name:	James J. Barlow
Email:	jbarlow@ncsa.uiuc.edu
Phone:	(217) 244-6403
Institution:	University of Illinois Urbana-Champaign
Department:	N.C.S.A.

Site	Login ID
ANL	jbarlow
Caltech	jbarlow
IU	tg-jbarlow
NCSA	jbarlow
ORNL	jbarlow
PSC	jbarlow
Purdue	jbarlow
SDSC	ux454965
TG	jbarlow
University of Texas at Austin	tg457851

	DNs
1.	/C=US/O=National Center for Supercomputing Applications/OU=People/CN=James J. Barlow
2.	/C=US/O=National Center for Supercomputing Applications/CN=James J. Barlow
3.	/C=US/O=Pittsburgh Supercomputing Center/OU=PSC Kerberos Certification Authority/CN=jbarlow/UID=jbarlow/emailAddress=jbarlow@PSC.EDU
4.	/C=US/O=Pittsburgh Supercomputing Center/OU=PSC Kerberos Certification Authority/CN=jbarlow/USERID=jbarlow/Email=jbarlow@PSC.EDU

TeraGrid IR Solutions

- **SSH Authentication database**
  - **Keeps record of all valid authentications**
  - **Checks for user logins from “new” sites**

```
Number of successful authentications: 34067

Number of unique remote IP addresses: 722
Top 15 remote IP addresses:
…

Number of different local hosts logged into: 187
Top 15 local hosts:
…

Number of different users who authenticated: 389
Top 15 users:
…

Different authentication types used: 7
   18842: publickey
   12779: hostbased
   1295: gssapi-with-mic
   1094: password
   44: keyboard-interactive/pam
   9: gssapi
   4: keyboard-interactive/cryptocard

Number of different ASN's: 143
Top 15 ASN's (total is the number of unique IP's within that ASN):
   284: 1224 - NCSA-AS - National Center for Supercomputing Applications
   63: 38 - UIUC - University of Illinois
   29: 7132 - SBIS-AS - AT&T Internet Services
…
```

TeraGrid IR Solutions

- Risk Analysis
  - Try to do yearly
  - Have done two so far (FRAP, NIST 800-26)

- Looking into user command profiling via process accounting logs
  - “A first step toward detecting ssh identity theft in hpc cluster environments: Discriminating masqueraders based on command behavior” by William Yurcik and Chao Liu
  - “Detecting SSH identity theft in HPC cluster environments using Self-organizing maps” by Claes Leufven

Questions?

jbarlow@ncsa.uiuc.edu

http://www.ncsa.uiuc.edu/~jbarlow/