

# Survey of Supercomputer Cluster Security Issues

George Markowsky

Linda Markowsky

Computer Science Department

University of Maine

# Outline

- Motivation
  - The Stakkato Intrusions
- Some Questions
- The Survey
  - Pool of Potential Participants
  - Responses
- Other Resources
- Recommendations
- Future Work

# Motivation

- Having done some teaching in cybersecurity and also having worked with supercomputer clusters, we became interested in the issue of supercomputer cluster security
- While there are many papers related to this topic, it was hard to find any status report on what is happening in this area

# Motivation

- We became curious about whether the security problems for supercomputer clusters were different from those faced by desktops
- Some of the reasons for suspecting that there might be differences are
  - Different types of data
  - Interest on the part of governments
  - Access to great computing power
  - Sophistication of users
  - Sophistication of attackers

# Motivation

- Additionally, it is clear that supercomputer cluster operators are aware of each other, so the supercomputer cluster network might be a tempting target
- Perhaps can help establish the state of the art and help cluster operators secure their clusters

**Breaking News»**

**Britain's Foreign Office confirms to CNN that five men killed today are Britons.**

**SEARCH**



The Web



CNN.com

**Search**

Home Page

World

U.S.

Weather

Business at CNNMONEY

Sports at SI.com

Politics

Law

**Technology**

Science & Space

Health

Entertainment

Travel

Education

Special Reports

Autos with EDMUNDS.com

**CRYOCARE**  
ProstateCancer.com

**CLICK HERE**  
For Information

SERVICES

Video

E-mail Newsletters

# TECHNOLOGY

## Report: Hacker infiltrated government computers

U.S. military installations, laboratories, and NASA hit last year

Tuesday, May 10, 2005 Posted: 2:22 PM EDT (1822 GMT)

WASHINGTON (CNN) -- The FBI confirmed Tuesday the accuracy of a New York Times report that software on routers, computers that control the Internet, were compromised last year by a hacker who claimed that he had infiltrated systems serving U.S. military installations, research laboratories, and NASA.



[advertiser links](#) [what's this?](#)



A Mind is a Terrible Thing to Waste

Visit [uncf.org](http://uncf.org)



[Read today's paper](#) · [Jobs](#)

Search:

Guardian Unlimited  Web

**Guardian**Unlimited **Special reports**

<a href="#">Home</a>	<a href="#">UK</a>	<a href="#">Business</a>	<a href="#">Audio</a>	<a href="#">Podcasts</a>	<a href="#">The Wrap</a>	<a href="#">News blog</a>	<a href="#">Talk</a>	<a href="#">Search</a>
<a href="#">The Guardian</a>	<a href="#">World</a>	<a href="#">News guide</a>	<a href="#">Arts</a>	<a href="#">Special reports</a>	<a href="#">Columnists</a>	<a href="#">Technology</a>	<a href="#">Help</a>	<a href="#">Quiz</a>

**Special report**  
United States of America

## Hacking trail leads to Swedish teen

Julian Borger in Washington  
Wednesday May 11, 2005  
[The Guardian](#)

Search this site

A Swedish teenager is being questioned over a daring internet attack that penetrated thousands of computer systems in the US, including military and Nasa websites, the FBI said yesterday.

Go to ...

[Special report: United States of America](#)

[United States of America archived articles](#)

A search is also under way in Britain and elsewhere in Europe for possible accomplices, a bureau spokesman said yesterday, but gave no further details.

"We've been working on this very closely with our international partners in Sweden, Britain, and others, and the criminal activity has stopped," the spokesman said.

# The Stakkato Intrusions: What Happened And What Have We Learned?

- Nixon, Leif (National Supercomputer Centre, Linkoping University)
  - Source: 6th IEEE International Symposium on Cluster Computing and the Grid, 2006. CCGRID 06, May 16-19 2006
- *During 15 months, from late 2003 until early 2005, hundreds of supercomputing sites, universities and companies worldwide were hit in a series of intrusions, with the perpetrator leapfrogging from site to site using harvested ssh passwords.*

# The Stakkato Intrusions: What Happened And What Have We Learned?

- The damage has been estimated to exceed \$100 million in the United States alone.
- These are known as the Stakkato intrusions.
- Nixon's talk covers case studies of performed intrusions, an analysis of why Stakkato could be so successful, and the story of how the suspect was finally tracked down and caught.

## The broader view



Some of the sites/companies/providers known to have been at the receiving end of an attack:

berkeley.edu	gatech.edu	rr.com	ucsc.edu
bonet.se	iastate.edu	rutgers.edu	ucsd.edu
brandeis.edu	jhu.edu	sdsc.edu	uiuc.edu
bredbandsbolaget.se	ki.se	seagull.net	umea.se
brown.edu	kralovopolska.cz	simons-rock.edu	umearc.se
bu.edu	kth.se	skanova.com	umn.edu
cam.ac.uk	liu.se	skogsbrynet.se	umu.se
cern.ch	liv.ac.uk	songnetworks.se	unige.ch
chalmers.se	lu.se	stanford.edu	uta.fi
cisco.com	mit.edu	technion.ac.il	utk.edu
columbia.edu	naqua.se	telia.com	uu.se
csbnet.se	nasa.gov	uchicago.edu	wsmr.army.mil
desy.de	nikhef.nl	uci.edu	
epfl.ch	pitt.edu	ucolorado.edu	

This is just a small sample; from August 2003 through March 2005 something like a thousand sites were attacked.

•Details from <http://www.nsc.liu.se/~nixon/stakkato.pdf>

# Some Questions

- What is the level of computer security expertise shown among cluster operators?
- To what extent are clusters targeted by organizations rather than random hackers?
- How common are physical or social engineering attacks?
- How sophisticated are the attackers?

# The Survey

- Some design considerations
  - Must not be intrusive
  - Must develop a relationship so people can trust us
  - Must not reveal weak spots to potential adversaries
    - Not security through obscurity
  - Must be short
  - Must preserve anonymity

# Pool of Potential Participants

- Put together a list for direct e-mail – well over 250 sent (early records incomplete)
- Posted questionnaire on various user groups
- Primary groups
  - USENET groups
  - beowulf.org mailing list
  - [www.securityfocus.com](http://www.securityfocus.com) e-mail lists

# Mail Was Sent To At Least 247 Different Addresses

<b>edu</b>	<b>159</b>	<b>nl</b>	<b>1</b>	<b>net</b>	<b>1</b>
<b>de</b>	<b>8</b>	<b>ru</b>	<b>2</b>	<b>ua</b>	<b>1</b>
<b>tr</b>	<b>1</b>	<b>ch</b>	<b>1</b>	<b>mil</b>	<b>4</b>
<b>gov</b>	<b>21</b>	<b>org</b>	<b>4</b>	<b>gr</b>	<b>1</b>
<b>au</b>	<b>13</b>	<b>hk</b>	<b>1</b>	<b>ie</b>	<b>1</b>
<b>ca</b>	<b>17</b>	<b>dk</b>	<b>3</b>	<b>us</b>	<b>4</b>
<b>be</b>	<b>1</b>	<b>uk</b>	<b>2</b>	<b>com</b>	<b>1</b>

# Responses

- We were very pleased by the response that we got to our efforts
  - 75+ surveys were completed.
  - Some consisted completely of No Answer, so only 75 surveys were used (earlier version of paper based on 61)
  - 9 letters were received
- No tracking of source of response
- We hope that once we post a more complete set of results that we will get additional feedback for future research

# **Question 1 : How frequently are your supercomputer clusters attacked relative to any desktops that might be in your laboratories?**


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 1 : How frequently are your supercomputer clusters attacked relative to any desktops that might be in your laboratories?



**NA = No Answer**  
**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

**Question 2 : How sophisticated are the attacks against your clusters compared to the attacks against any desktops that might be in your laboratories?**

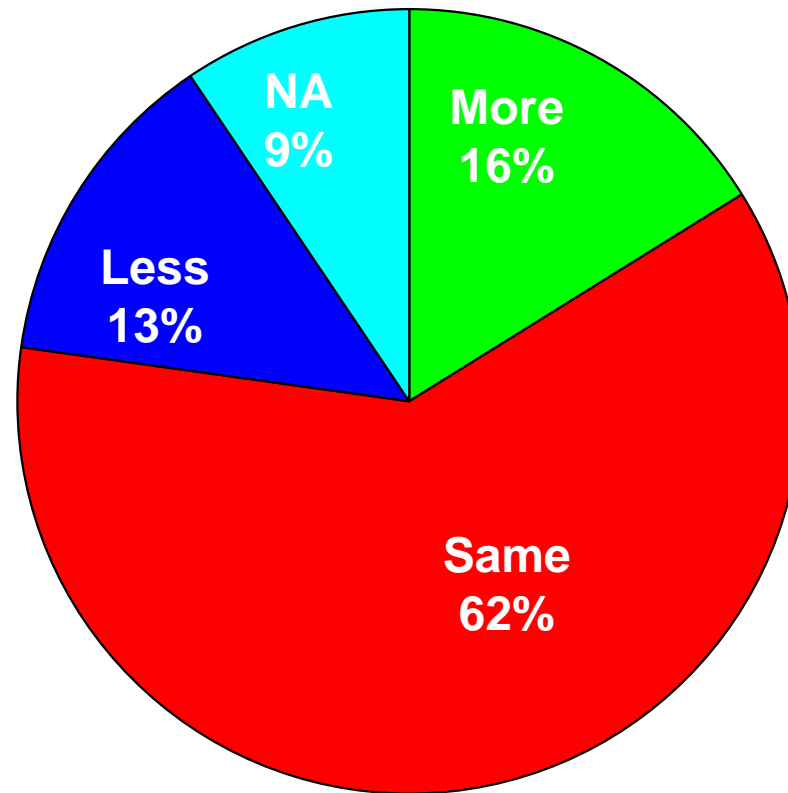
**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 2 : How sophisticated are the attacks against your clusters compared to the attacks against any desktops that might be in your laboratories?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# **Question 3 : Are there any IP addresses that regularly try to break into your clusters?**


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 3 : Are there any IP addresses that regularly try to break into your clusters?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# **Question 4 : Has anyone ever tried a man-in-the-middle type of attack against any of your clusters?**


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 4 : Has anyone ever tried a man-in-the-middle type of attack against any of your clusters?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 5 : Have you ever been attacked from foreign IP addresses?


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 5 : Have you ever been attacked from foreign IP addresses?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 6 : Have your clusters ever been attacked by foreign interests?


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 6 : Have your clusters ever been attacked by foreign interests?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 6 Comments

- How can people tell one way or the other
- We were surprised to find that roughly half the respondents had a definite opinion on this matter – we were expecting that the overwhelming majority would be in the Not Sure camp
- This suggests lots of interesting follow-up questions

# **Question 7 : Has anyone ever tried a physical approach to either disrupt a computation or to steal data?**


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 7 : Has anyone ever tried a physical approach to either disrupt a computation or to steal data?



NA = No Answer

75 Responses

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

**Question 8 : Has anyone ever tried to bribe or otherwise co-opt one of the cluster staff into helping with compromising the security?**


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 8 : Has anyone ever tried to bribe or otherwise co-opt one of the cluster staff into helping with compromising the security?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Comments on Questions 7 & 8

- The positive responses in questions 7 (7 yes) and 8 (6 yes) were generally different organizations (11 different respondents)
- It would seem worthwhile to understand this better
- It would seem that about 15% of the respondents experienced either a physical approach or an attempt to co-opt a member of the staff

**Question 9 : How many times has security been breached on one of your supercomputer clusters over the past three years that resulted in either downtime or lost data?**

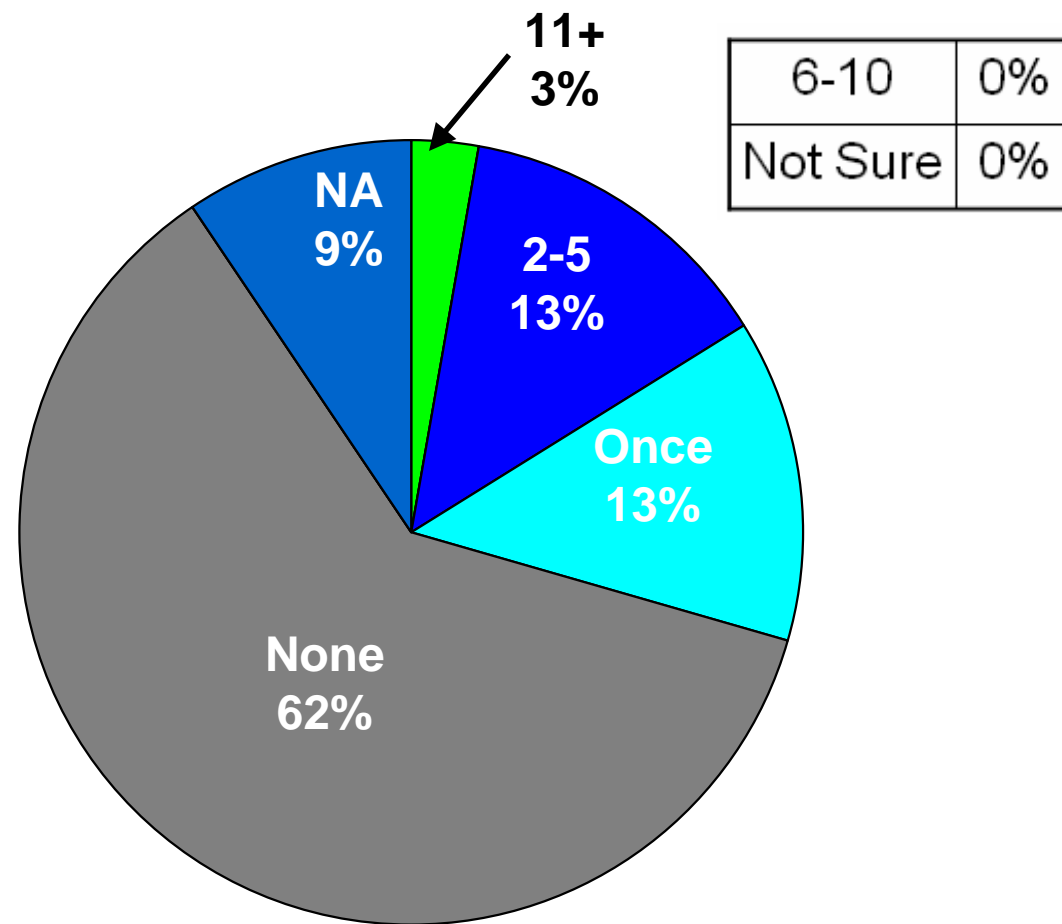
**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 9 : How many times has security been breached on one of your supercomputer clusters over the past three years that resulted in either downtime or lost data?



**NA = No Answer**  
**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# **Question 10 : Does your center have a person whose primary responsibility is cluster security?**


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 10 : Does your center have a person whose primary responsibility is cluster security?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 11 : Do you run an intrusion detection system on your clusters?


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 11 : Do you run an intrusion detection system on your clusters?



**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 11 Comments

- It seems that less than 50% of the respondents stated that they had an intrusion detection system running on all their systems
- It would be a good idea for more cluster operators to run intrusion detection software on all their clusters

# Question 12 : How often do you check for rootkits?


**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 12 : How often do you check for rootkits?



NA = No Answer  
75 Responses

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 12 Comments

- It would seem worthwhile for more cluster operators to check for rootkits on a daily basis
- The Stakkato Intrusions were characterized by the intruder installing rootkits on many servers

# Question 13 : How often do you run backups on your clusters?

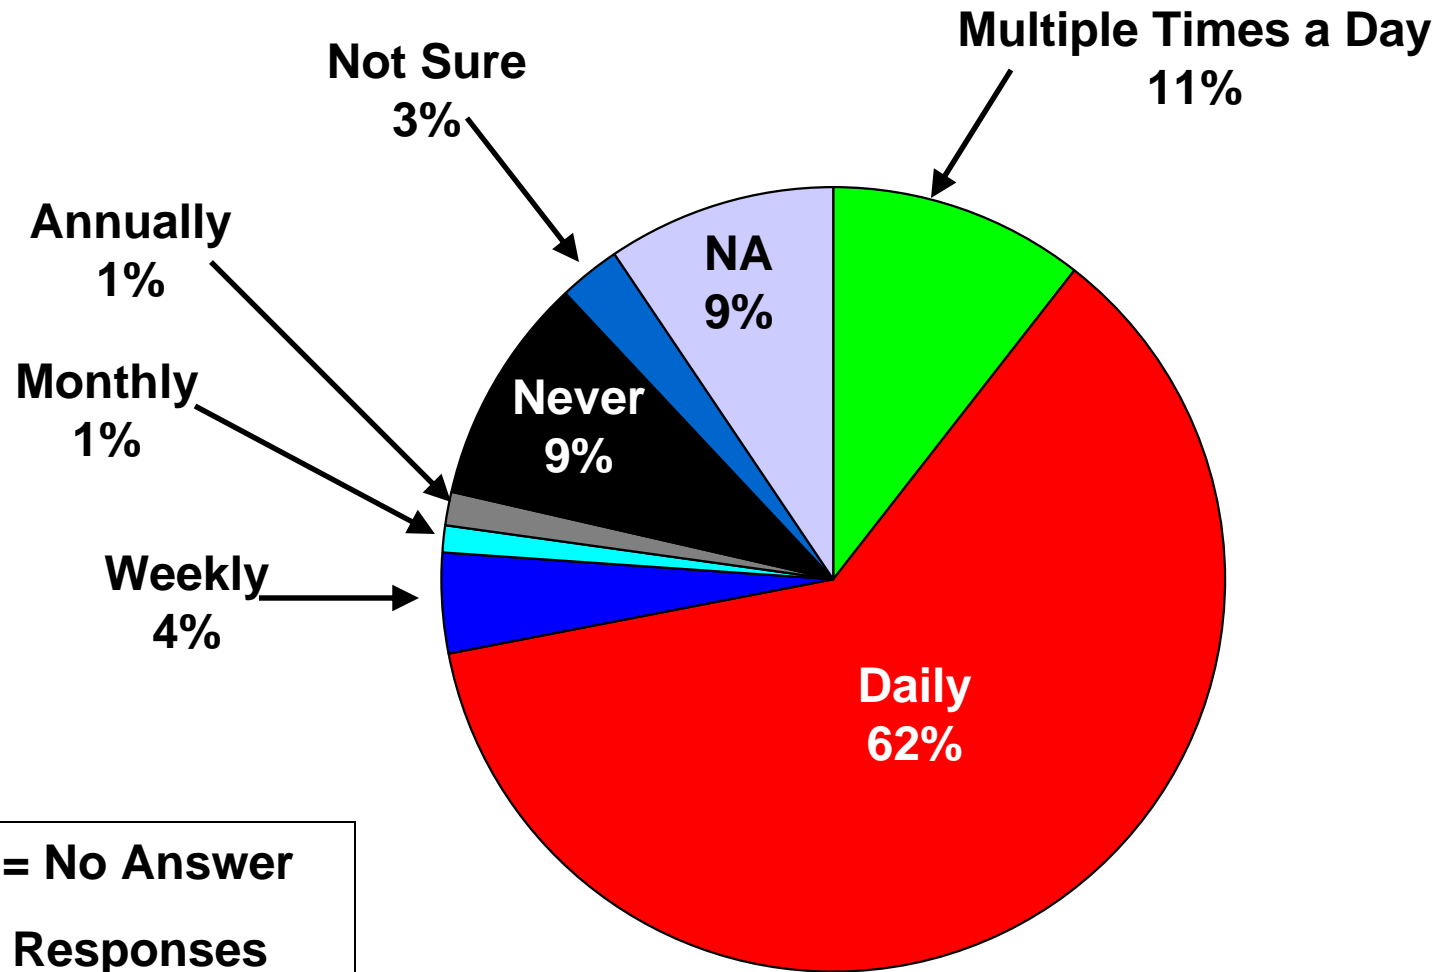
**NA = No Answer**

**75 Responses**

June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Question 13 : How often do you run backups on your clusters?



June 23, 2007

Markowsky & Markowsky Supercomputer Cluster Survey

# Some Letters

Type of Letter	Number
Verification	2
Note of Duplication & Support	1
Interest in Results	2
Comments & Suggestions	4

# Letter 1

- *I believe that the most important question is missing:*
- *"Have you classified or proprietary data on your cluster?"*
- *Otherwise the cluster is "not worth attacking". A normal university or hobby cluster will be defended differently than a sensitive one. Of course the question itself is sensitive.*

# Response to Letter 1

- We were reluctant to ask questions of this type because we did not want people to think that we were footprinting their sites
- We are thinking about how to best deal with this issue in the future
- We feel that all clusters are worth attacking because they can provide a "trusted" base for further attacks as was done in the Stakkato Intrusions

# Response to Letter 1

- All clusters are worth attacking since they can provide a lot of computational power for such activities as password cracking

# Letter 2

- *I would like to point out that the use model for clusters is what drives most of the important security risks and is the primary threat.*
- *Laptop/desktop and standard commercial activities do not normally involve giving a large number of users shell access. Very few HPC systems run through portals where the OS is isolated from the user.*

# Letter 2

- *So the starting point for a cluster is where the 85% of the standard security stops (don't let the bad guy get a shell).*
- *Compromised user desktops will immediately lead to a user level compromise of the cluster. If you don't manage the user desktops (they might not be yours) you must assume that there is an account compromise.*
- *The important question is what do you do then.*

# Response to Letter 2

- The problem of protecting internal resources from compromised user desktops/laptops is very serious – it has led some organizations to put up firewalls between resources and users!
- We believe that a variety of steps can be taken to secure clusters even in case one or more user computers get compromised

# Response to Letter 2

- Even if a keylogger has been installed on the desktop, there are ways of protecting the cluster, e.g.,
  - Users could be required to use key-based authentication, with the key kept on removable media, not on the hard drive, or
  - Systems can use keys together with one-time passwords
- User-level compromise is still possible, but it would hardly be "immediate".

# Letter 3

- *I think that the best way to keep clusters secure is to hide the compute nodes on a private network and keep them totally inaccessible from the outside world. Any access to compute nodes is through the master node.*
- *Compute nodes, if they need to be installed or updated, do so from a repository on the master node. The master node, then, is firewalled on the public interface, with only a few ports open-- ssh, http, https, and so on.*

# Letter 3

- *The internal interface is wide open, since you never know what researchers are going to want to run, but even if they start up something that tries to listen on the outside interface it remains unreachable by naughty people.*

# Letter 3

- *Once you lock it down like that, the number of potential problems decreases pretty drastically. For some things (apache server-status and so on), I have access restricted to my workstation in the apache config. For compute clusters, I can't see the need for lots of off-campus access, so even locking things down to on-campus users doesn't seem overly restrictive.*

# Letter 3

- *Overall, I've found our compute resources to be much less likely to attract security problems than our public systems. We aren't doing any DoD research or anything, though.*

# Response to Letter 3

- We feel that it enhances security by having defenses installed on compute nodes as well as the master node
- Many users are security naive and we think that it is good to have some controls over what they run

# Letter 4

- Very long and very interesting letter – one tidbit:
- *Over the decades I have accumulated a number of interesting anecdotes on security -- major cracks of systems in our medical center, a Bulgarian grad student who engaged in rampant data theft on a chemistry cluster shipping research data on rational drug design back to Bulgaria, and more.*

# Letter 4

- *The worst incidents that resulted in actual damage were of this sort -- "inside jobs" where lax security boundaries INSIDE an organization permitted unauthorized access, at least until they were detected and bopped.*

# Response to Letter 4

- This letter also suggested collecting some war stories
- It might be good to have such a collection that could be accessed by interested people

# Other Resources

- Uncovered many relevant papers and resources
- We will put materials and links at

<http://www.cs.umaine.edu/~markov/clustersecurity>

- We will give a quick look at some of the papers that we think would be of general interest by quoting excerpts from their abstracts

# Searching For Open Windows And Unlocked Doors: Port Scanning In Large-scale Commodity Clusters

- by Lee, Koenig, Meng, and Yurcik
- *... By correlating the open network ports observed on cluster nodes with other emergent properties - such as active processes and the contents of important system files security analysts can make insightful observations that can greatly restrict the actions that an attacker can carry out undetected.*

# Intrusion-tolerant Server Architecture For Survivable Services

- by Min
- .. *deliver intended services transparently to the clients even when a computing node fails due to failures, intrusions, and other threats ... deliver only secure results to the client.*

# Nvisioncc: A Visualization Framework For High Performance Cluster Security

- by Yurcik, Meng, and Kiyancilar
- *... a framework for effective visualization of a high performance cluster security ... GUI screenshots from a security visualization tool based on this framework ...*

# Detecting Anomalies In High-Performance Parallel Programs

- by Florez, Liu, Bridges, Vaughn, and Skjellum
- *... a number of different types of irregularities can occur including those that result from intrusions, user misbehavior, corrupted data, deadlocks or failure of cluster components ... artificial intelligence (AI) techniques ... can be used to implement a lightweight monitoring and detection system for parallel applications on a cluster of Linux workstations ... monitoring of MPI programs can be achieved with high accuracy and in some cases with a 0% false positive rate in real-time ... the added computational load on each node is small.*

# The Trellis Security Infrastructure For Overlay Metacomputers And Bridged Distributed File Systems

- by Lu, Closson, Macdonell, Nalos, Ngo, Kan, and Lee
- *Researchers often have non-privileged access to a variety of high-performance computer (HPC) systems in different administrative domains, possibly across a wide-area network. Consequently, the security infrastructure becomes an important component of an overlay metacomputer: a user-level aggregation of HPC systems. The Trellis security infrastructure (TSI) is layered on top of the widely-deployed secure shell (SSH) and systems administrators only need to provide unprivileged accounts to the users.*

# The Trellis Security Infrastructure For Overlay Metacomputers And Bridged Distributed File Systems

- *The contribution of TSI is in demonstrating that a single sign-on (SSO) system, for a variety of use-case scenarios, can be implemented without requiring a completely new security infrastructure.*

# Survival of Internet Applications: A Cluster Recovery Model

- by Aung, Park, and Park
- *... Prevention is a necessary but not a sufficient component of disaster. In this case, we have to prepare thoroughly for reducing the recovery time and get the users back to work faster.*

# NIDS Architecture for Clusters

- by Gadaud
- *... a NIDS architecture which addresses the same constraints as a cluster: efficiency, scalability and reliability. It is based on the cluster paradigm.*

# Instant Attack Stopper in Infiniband Architecture

- by Lee, Kim, Yum, and Yousif
- *... a scheme, referred to as the Instant Attack Stopper (IAS) that can instantly confront security attacks in a cluster. Specifically we provide detailed implementation methods of IAS in InfiniBand Architecture (IBA) - a new promising communication standard for future System Area Networks (SANs) and clusters. IAS focuses on removing malicious communication on the IBA fabric among processes involved in an attack, which is accomplished through the proposed Security Management Agent (SeMA).*

# Recommendations

- A number of respondents seem to feel that things are under control
- The success of the Stakkato Intrusions suggest that no one should be overconfident
- We feel that all cluster operators will benefit from increased information exchange between cluster operators

# Future

- Future Questionnaires
- Website
  - References
  - Tools
- Additional Sessions Focusing on Security of Supercomputer Clusters

# Contact Information

- `markov@maine.edu`
- `http://www.cs.umaine.edu/~markov`