Monitoring-Driven Security and Dependability

Ravishankar K. Iyer
Phuong Cao, Cuong Pham Z. Kalbarczyk, Hui Lin, Homa Alemzadeh….

Coordinated Science Laboratory and
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
www.csl.edu/DEPEND

What Happens in an **Internet Minute?**

And **Future Growth** is **Staggering**

Today, the number of networked devices = the global population

By **2015**, the number of networked devices = **2x** the global population

In **2015**, it would take you **5 years** to view all video crossing IP networks each **second**

intel

SOURCE: http://www.intel.com/content/www/us/en/communications/internet-minute-infographic.html

Clouds : Growing Number of Outages

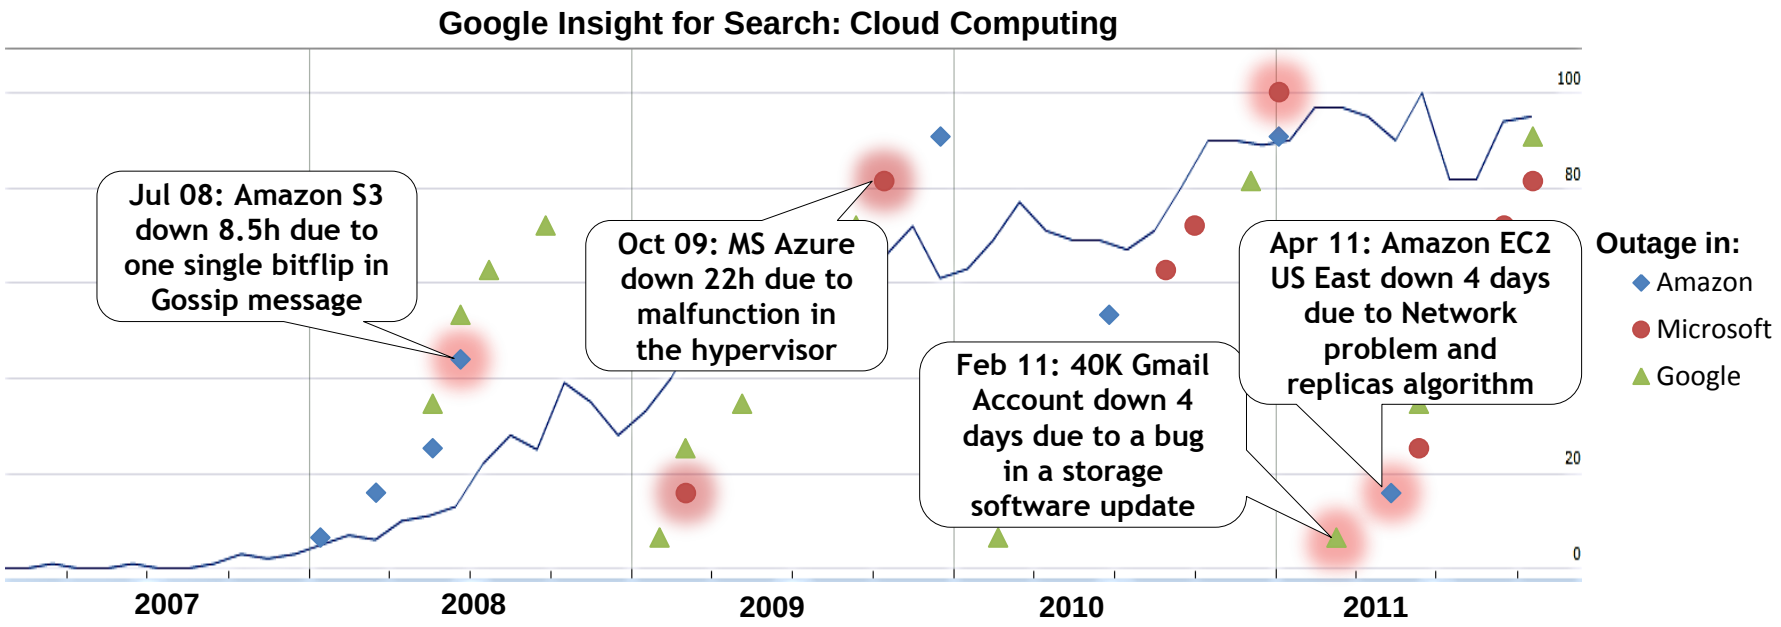

- Providing a higher level of reliability and availability is one of the biggest challenges of Cloud computing

From C. Pham in R. Iyer, Z. Kalbarczyk, and N. Nakka, "Dependable Computing: Design and Assessment," forthcoming text, book in 2013

Clouds: Security Problems

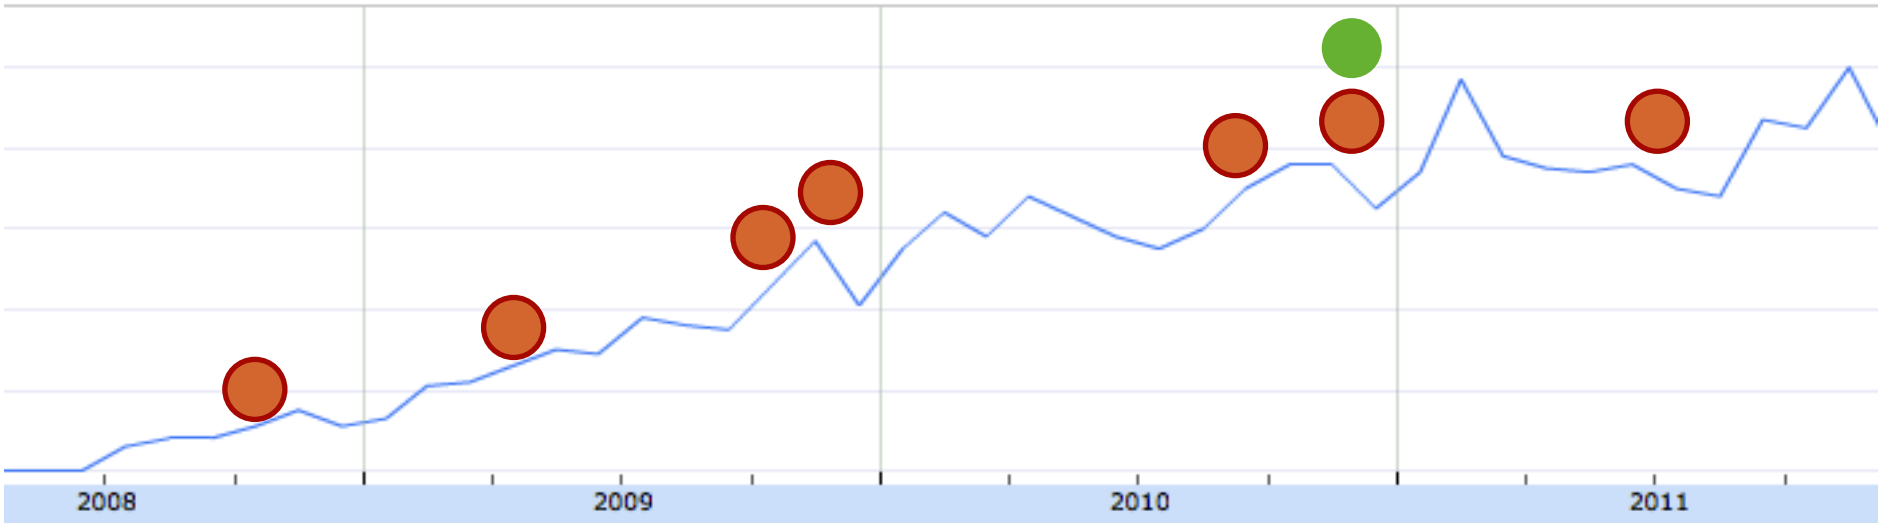

Jul'08 - Spammers set up mail spamming instances in the Amazon's EC2 cloud.

Sep'10 - Google Engineer Stalked Teens, Spied on Chats

Apr'09 - Texas datacenters operations are suspended for FBI investigation.

Dec'10 - Microsoft BPOS cloud service hit with data breach

Nov'09 - Side channel attack of Amazon's EC2 service.

June'11 - Dropbox: Authentication Bug Left Cloud Storage Accounts Wide Open

Dec'09 - Zeus crime-ware using Amazon's EC2 as command and control server.

Dec'10 - Anonymous hacker group failed to take down Amazon

Heath Care Example: Surgical Robot Accidents

A Comparison to Aviation Industry

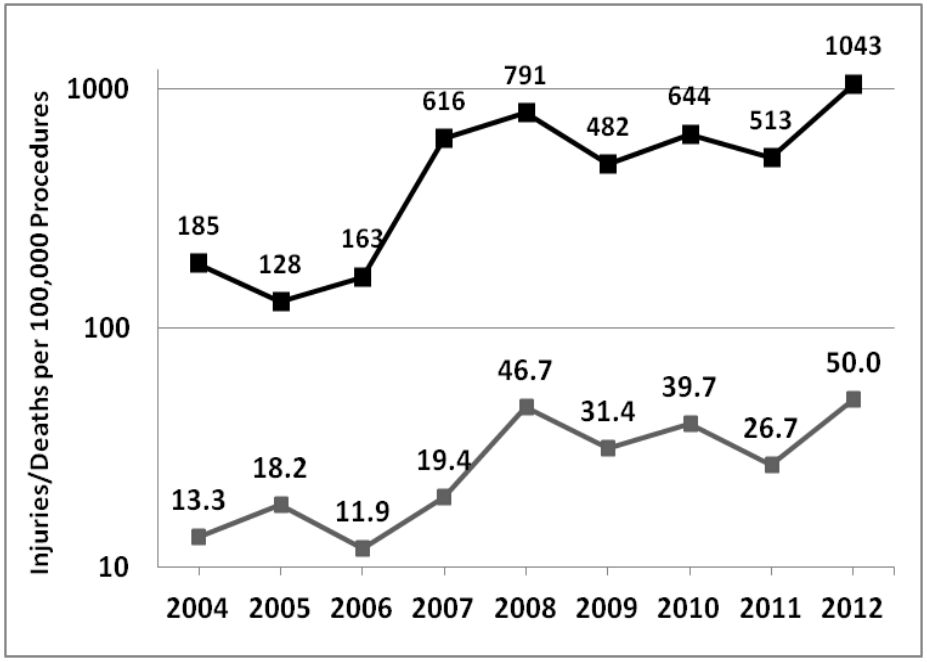

Rates of Robotic Surgery Injury or Death Reports (Robotic Surgery Accidents per 100,000 procedures)

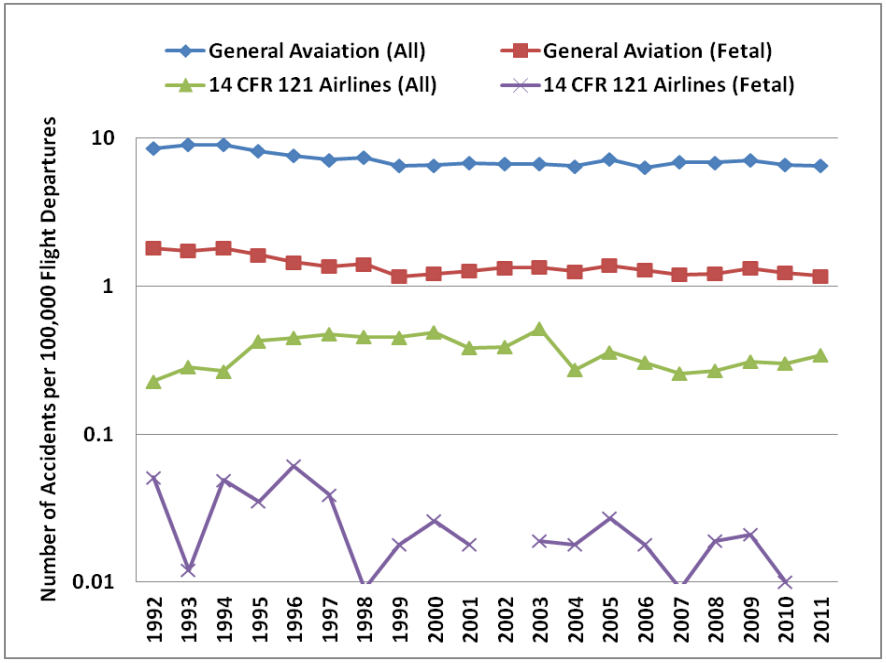

Aviation Accident Statistics by NTSB (1992-2011) (Accident per 100,000 flight departures)

From: http://www.ntsb.gov/data/aviation_stats.html

Robotic Surgery Accident Rates:
All: 128 to 1043
Safety-Critical: 11.9 - 50

1-2 order of magnitudes higher accident rates

Aviation Accident Rates:
All: 0.23 to 8.51
Fatal: 0.01 to 1.81

H. Alemzadeh, R. K. Iyer, J. Raman, "Safety Implications of the Robotic Surgery: Analysis of Recalls and Adverse Events of da Vinci Surgical System", DEPEND Technical Report, June 2013.

Major Challenges

- Develop and enforce security and reliability policies?
- Continuously monitor (and respond to) attacks and failures
- Assured virtual environments
- Estimate, Validate Benchmark

Measurement Driven Approach

Analysis of Security Meaasurements from a Large System: NCSA Case Study

- Goals:
  - Provide the system-level characterization of incidents and evaluate the intricacies of carrying out successful attacks
  - Design attack independent protection strategies to reduce the number of missed incidents and false positives
  - Demonstrate techniques in an experimental testbed
- Challenges

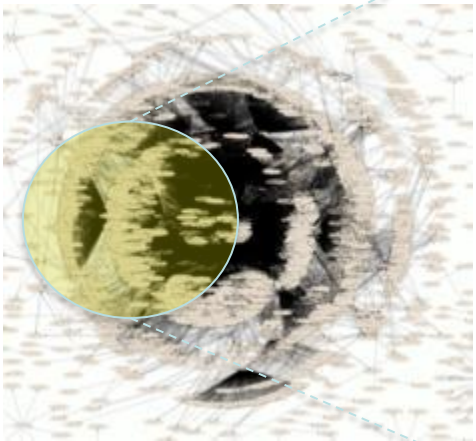

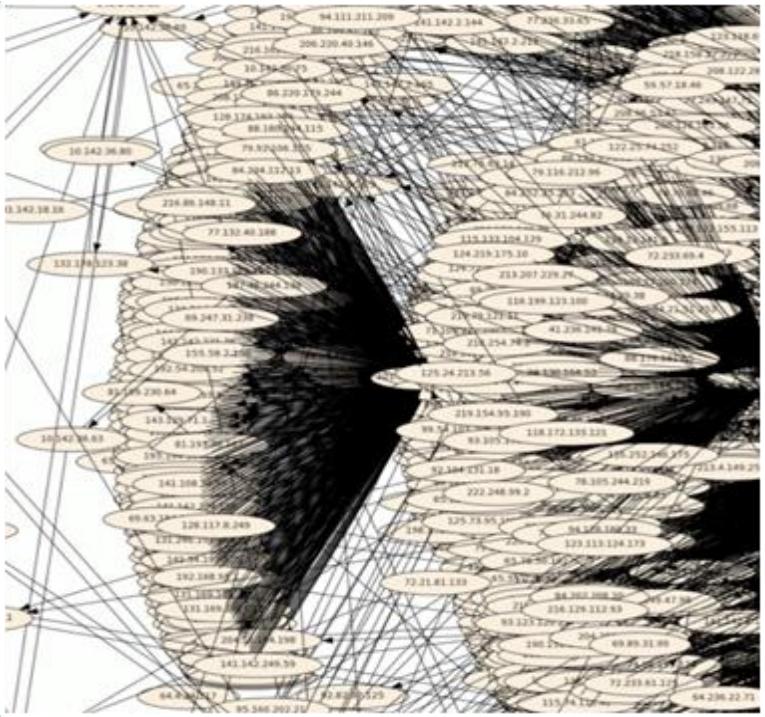

Five-Minute Snapshot of In-and-Out Traffic within NCSA

Five-Minute Snapshot of In-and-Out Traffic within NCSA

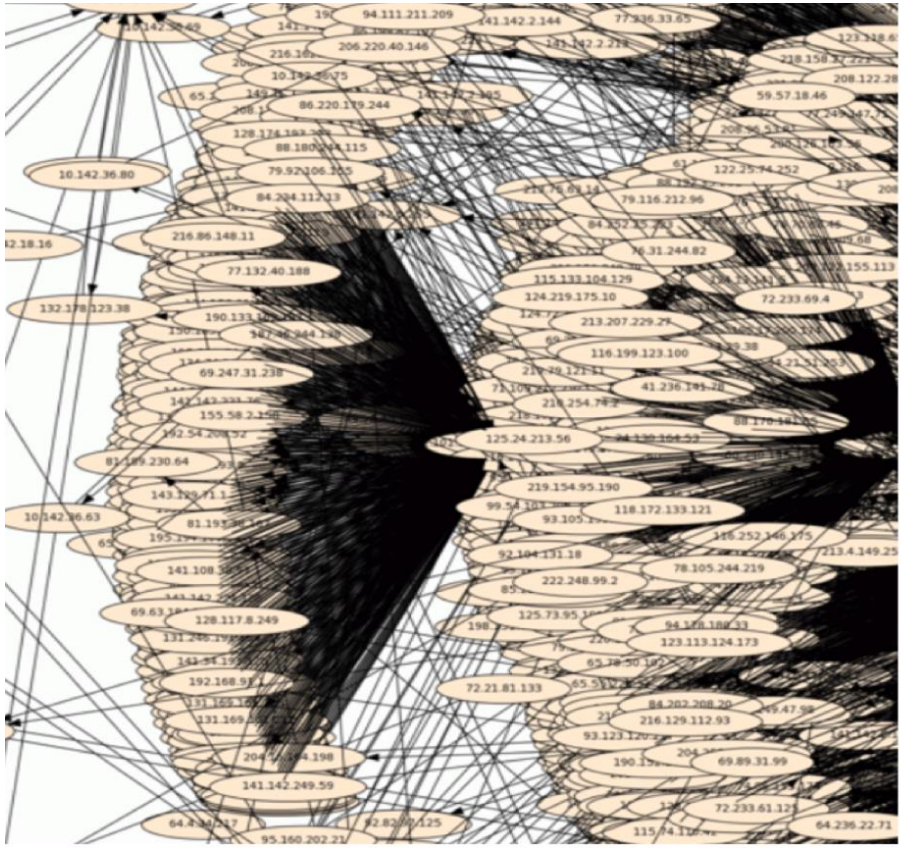

(a)

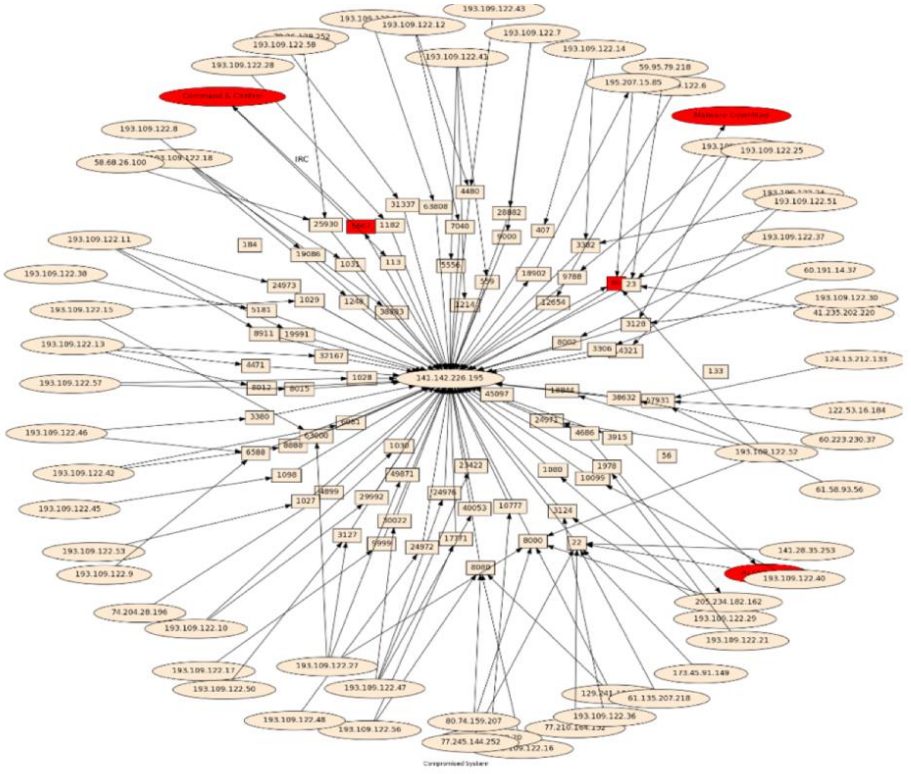

(b)

Approach

- Analysis of data logs on security incidents at National Center for Supercomputing Applications (NCSA)
  - Over 5000 machines accessible across the world
  - Total number of investigations: **212**
  - Real incidents: **178**
  - False positives: **34**

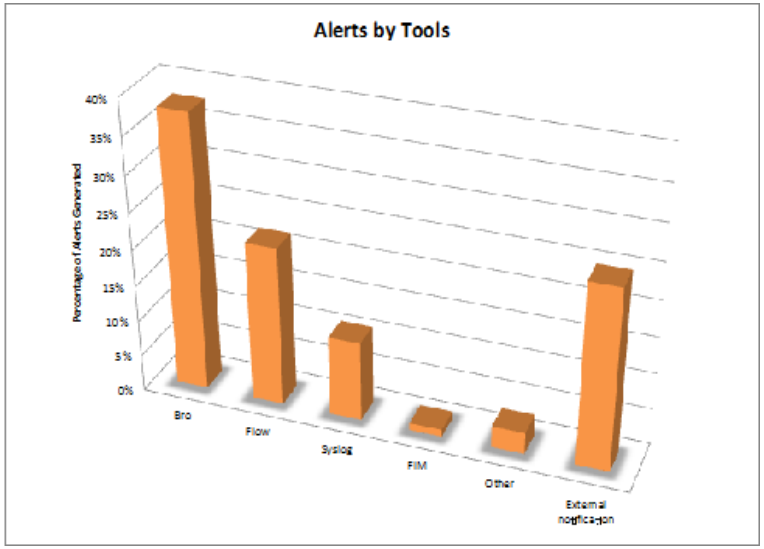

- Monitoring Tools
  - **Bro IDS:** performs deep packet inspection of network traffic
  - **Network flows:** monitored using Argus and nfdump
  - **Syslog:** Simple Event Correlation engine (SEC) generates alerts based on rule sets
  - **File Integrity Monitor (FIM):** alerts on changes to critical system files

Sample Results: Credentials Stealing Incidents

- Initial investigation of security incidents indicated that nearly 26% (32/124) of the incidents analyzed involved credentials stealing
- 31 out of 32 incidents attackers came into the system with a valid credential of an NCSA user account
  - Attackers rely on their access to an external repository of valid credentials to harvest more credentials
  - Availability of valid credentials makes boundary protections (e.g., reliance only on a firewall) insufficient for this type of attacks.
  - More scrutiny in monitoring user actions is required

Analysis of an Example Incident

(Credentials Stealing Category: Total 32 incidents)

- **An IDS alert shows suspicious download** on a production system (victim: *xx.yy.ww.zz*) using http protocol from remote host *aa.bb.cc.dd*.

May 16 03:32:36 %187538 start xx.yy.ww.zz:44619 > aa.bb.cc.dd:80
May 16 03:32:36 %187538 GET /.0/ptrat.c (200 "OK" [2286] server5.bad-host.com)

- The file is suspect because
  - This particular system is not expected to download any code apart from patches and system updates, and then only from authorized sources
  - The downloaded file is a C language source code
- The server the source was downloaded from not a formal software distribution repository.
- ***The alert does not reveal what caused the potentially illegal download request***

Correlations with Other Logs

- **Network flows reveal further connections with other hosts** in close time proximity to the occurrence of the download:
  - SSH connection from IP address 195.aa.bb.cc
  - Multiple FTP connections to ee.ff.gg.hh, pp.qq.rr.ss.

```
09-05-16 03:32:27 v tcp 195.aa.bb.cc.35213 -> xx.yy.ww.zz.22 80  96  8698 14159   FIN
09-05-16 03:33:36 v tcp xx.yy.ww.zz.44619 -> aa.bb.cc.dd.http 8  6  698 4159   FIN
09-05-16 03:34:37 v tcp xx.yy.ww.zz.53205 -> ee.ff.gg.hh.ftp 1699 2527 108920 359566 FIN
09-05-16 03:35:39 v tcp xx.yy.ww.zz.39837 -> pp.qq.rr.ss.ftp 236  364  15247  546947 FIN
```

- ***SSH connection record does not reveal***
  - ***Whether authentication was successful***
  - ***What credentials were used to authenticate the user***

Correlation with *syslog* Alerts

- *syslog* confirms a user login from *195.aa.bb.cc*, which is unusual, based on the known user profile and behavior patterns

---

May 16 03:32:27 host sshd[7419]: Accepted password for user from 195.aa.bb.cc port 35794 ssh2

- ***Four data points established from the analysis***
  - ***A suspicious source code was downloaded,***
  - ***The user login occurred at nearly the same time as the download,***
  - ***First time login from IP address 195.aa.bb.cc,***
  - ***Additional communication on other ports (FTP)***

Additional (Manual) Analysis

- Search of all files owned or created by this user found a footprint left behind by a credential-stealing exploit.

```
-rwxrwxr-x 1 user user 3945 May 16 03:37 /tmp/libno_ex.so.1.0
```

- ***The additional analysis showed***
  - ***The library file libno_ex.so.1.0 is known to be created when an exploit code for a known vulnerability (cve-2009-1185) is successfully executed***
  - ***File is owned by the user whose account was stolen and used to login to the system***
  - ***The attacker obtained root privileges in the system and replaced the SSHD daemon with a trojaned version***
    - ***Harvesting more user credentials***

Sample Results: Missed Incidents

Distribution of Incidents by Type 2004-2011

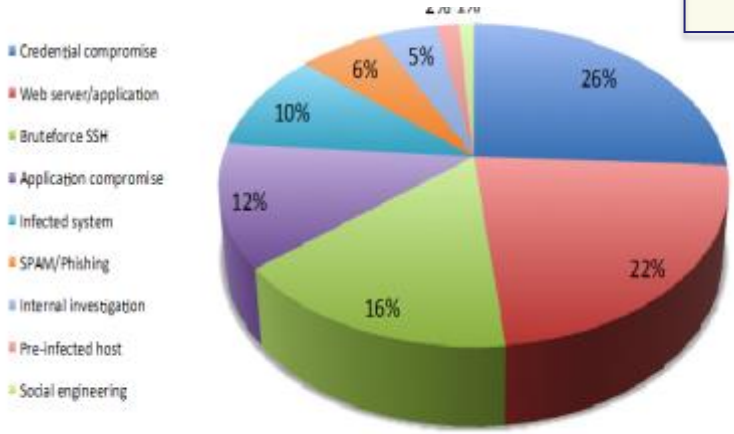

Severity of Incidents by Monitoring Tools 2004-2011

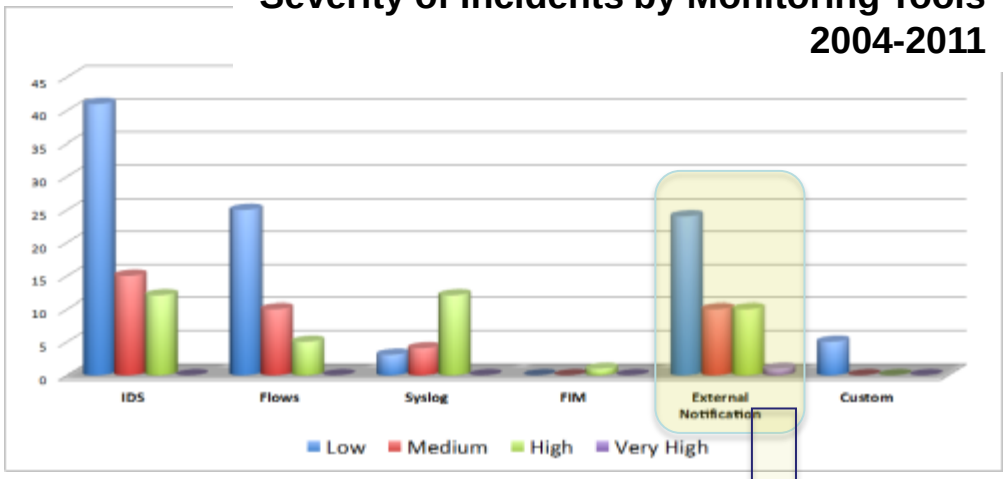

Cause of missed incidents	Examples	#
Increased sophistication in attacks	A peer site gets compromised and attacker logs-in with stolen credentials; zero-day exploits	6
Lack of signatures	Exploit of VNC null string authentication vulnerability	7
Admin misconfiguration	Web share world writable access or root login to accept any password	5
Inability to distinguish traffic anomalies in the network	Web defacement or use of web server to host malware; bot command and control traffic	10
Misconfiguration of security monitoring tools	Routers stop exporting the flows to central collector which prevents alerting	1
Inability to distinguish true positives from false positives	Human error	2
Inability to run monitors on all hosts and file systems due to cost	Limited deployment of file integrity monitors on non-critical systems	3

Significant portion of undetected incidents have high and very high impact (severity)

25% of incidents are missed (undetected)

Summary of Measurements

- Introduce data-driven methodology to evaluate detection capabilities of security monitoring system and characterize incidents
- No single available tool can perform the kind of analysis presented
- Need to correlate:
  - data from different monitors
  - system logs
  - human expertise
- Need to develop techniques to pre-empt an attacker actions
  - potentially let the attacker to progress under *probation* (or tight scrutiny) until the real intentions are clear

Key Findings

- Over half (57%) of incidents are detected by IDS-Bro (31%) and NetFlows (26%) monitors
- 27% of incidents are not detected by any alert
- 26% of the incidents involved credentials stealing
  - an attacker becomes an insider
- Nearly 39% of the incidents are detected in the last stage of the attack (attack-relay/misuse)
- Anomaly-based detectors are seven times more likely to capture an incident than are signature-based detectors
  - signatures are specialized to detect the presence (or download) of a known malicious binary but can be easily subverted

Identifying Compromised Users in Shared Computing Infrastructures: a Data-Driven Approach

Goals:

- protect integrity and confidentiality of data and applications from unauthorized and malicious access
- understanding characteristics of security alerts
- design automated approach to support incident investigation
- validating the approach against real incident data

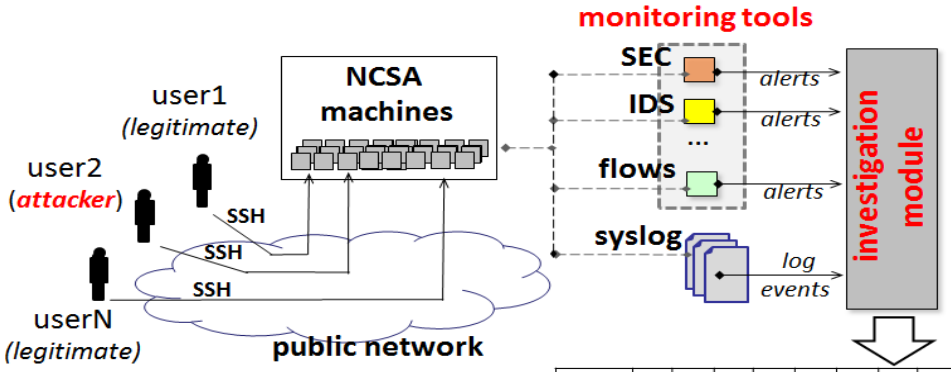

The module infers the users *potentially* responsible for the alerts observed in the security logs

Given a user, Ai is **1** if the user has been related to an alert of type Ai (user/alert table representation)

	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12	A13	A14	A15
user1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
user2	1	0	0	1	0	0	0	0	0	0	1	0	0	0	1
...	...														
userN	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0

Incident analysis using Bayesian network

user	1	0	0	1	0	0	0	0	0	0	1	0	0	0	0

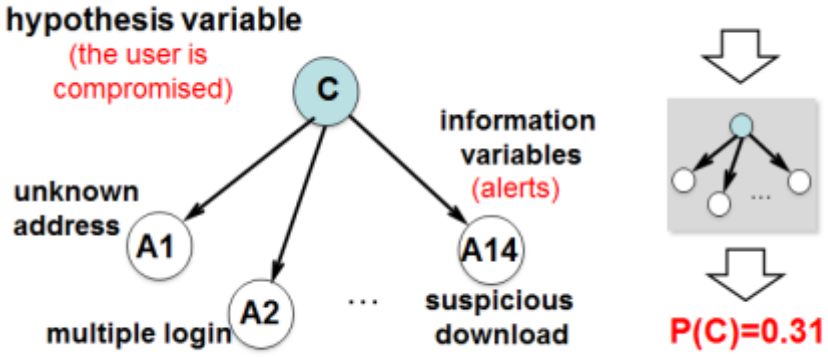

The **representation** produced by the investigation module is queried against the network. The output is ***the probability P(C), of the user being compromised.***

Sample Results

- Key findings:
  - it is feasible to define a classification threshold *to discriminate suspicious from compromised users*
  - classification conducted via the Bayesian network approach allows reducing the number of false compromise indications by about 80%
  - the network supports the investigation of new incidents

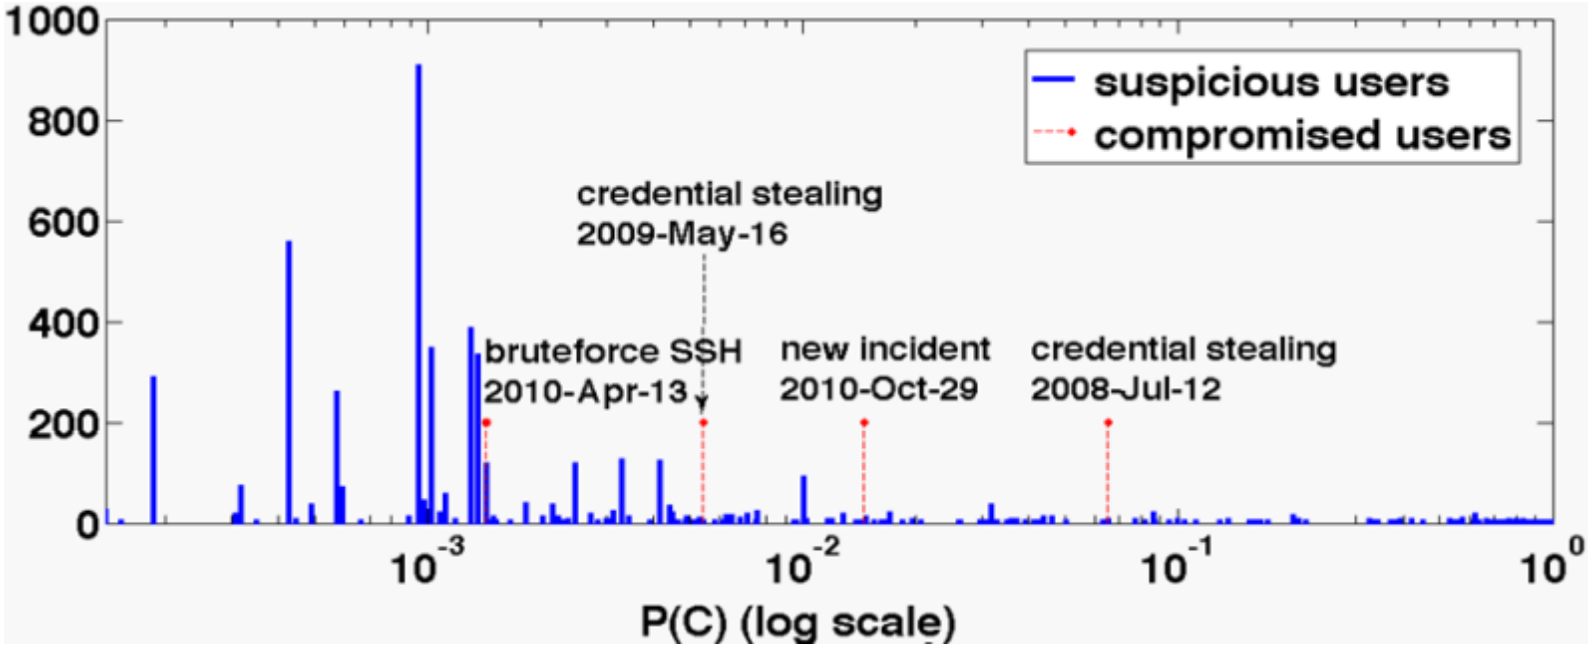

MONITORING DRIVEN TRUST

NCSA Operational Data in a day

Data-Driven Security

1000
users

5 millions
connections

10 millions
log lines

140
alerts

Operational Data

Learn attackers' behaviors

Preempt attacks

Block malicious actions

Fundamental Tradeoffs

How much does monitoring cost?

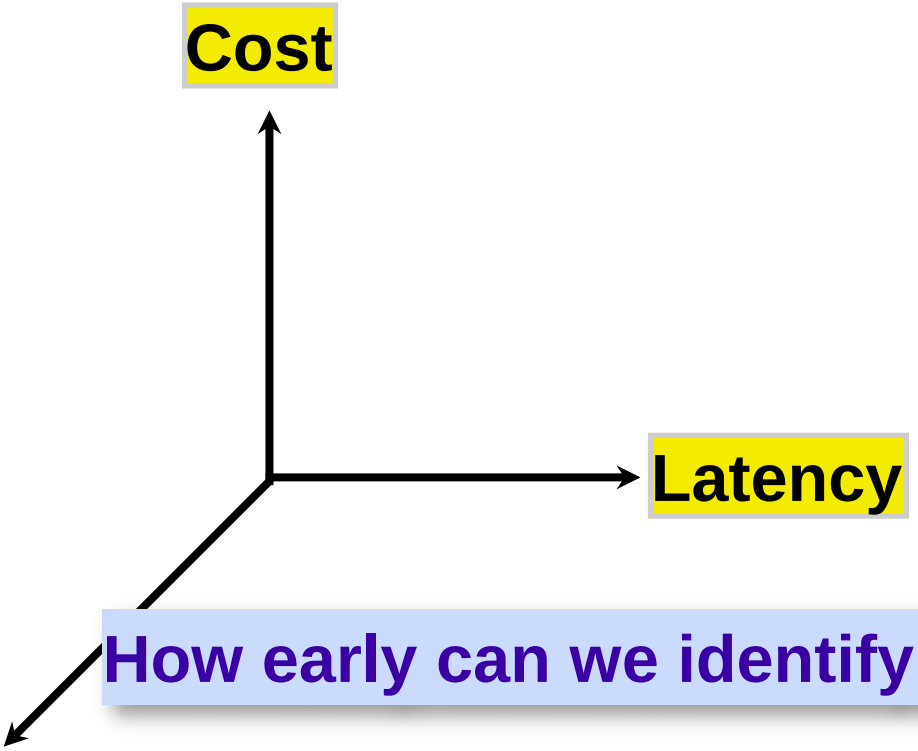

What is the desired detection accuracy?

Ecexution Under Probation

Real-time Analytics

Preempt attacks

Attack Prediction

Continuous Monitoring

Conclude and block attacks

Probation Environment

Learn attackers' behaviors

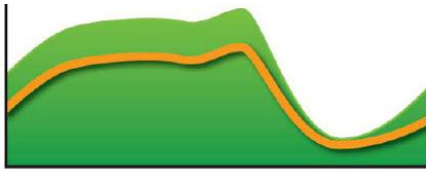

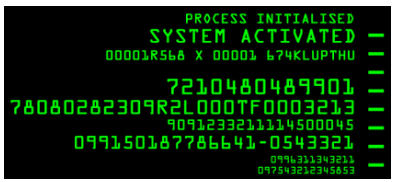

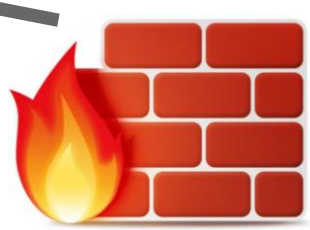

Execution Under Probation Effectiveness

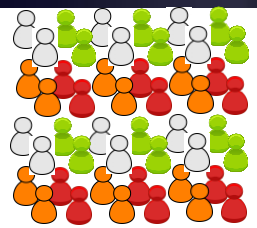

1021 users + alerts

Compute Suspicion Score

1. Compute Suspicion Score using:
Past: use ground truth data to compute likelihood
Present: use alert disorder, alert rate, and decay factor
2. Select top suspicious users

232 suspicious users

Monitor in Probation Environment

1. Look for users that generate more than three alerts in probation environment. They are potential attackers.
2. Return other users to normal execution environment.

42 had more than 3 alerts

Block Suspicious Activities

Block suspicious commands, e.g., “sudo” to prevent privilege escalation. We use a learned dictionary of suspicious commands.

14 attackers of a total 15 attackers

That means 90+% detection rate. We miss 6.67% of attacks - considerably better than 27% misdetection rate of previous study (Aashish et. al., DSN 2011)

Real-time Monitoring Dashboard

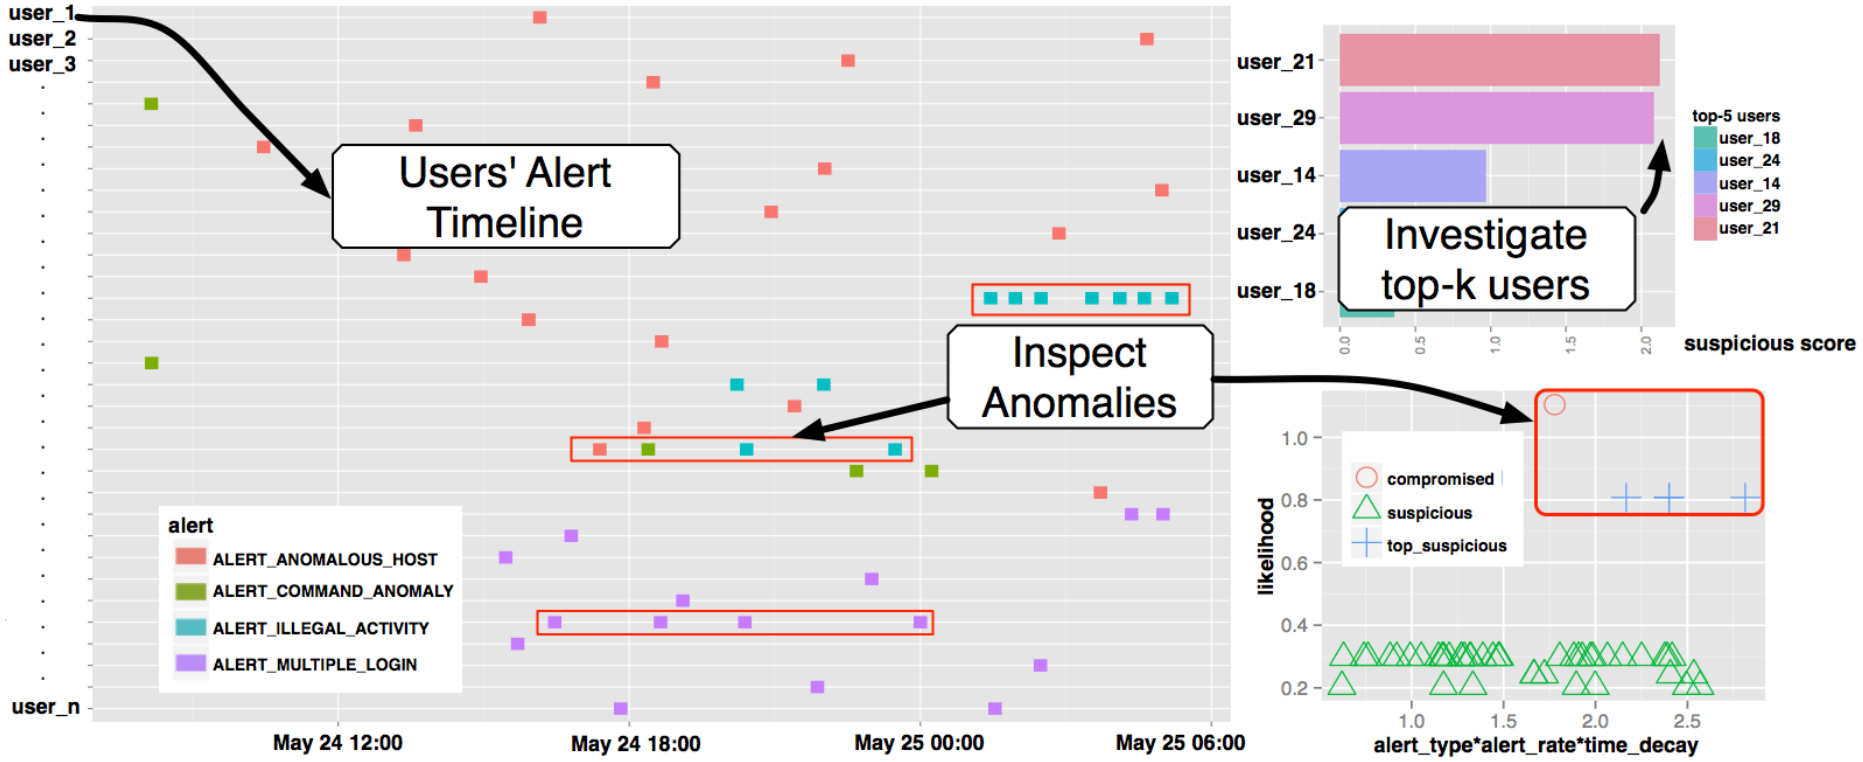

Power Grid Example: SCADA Testbed

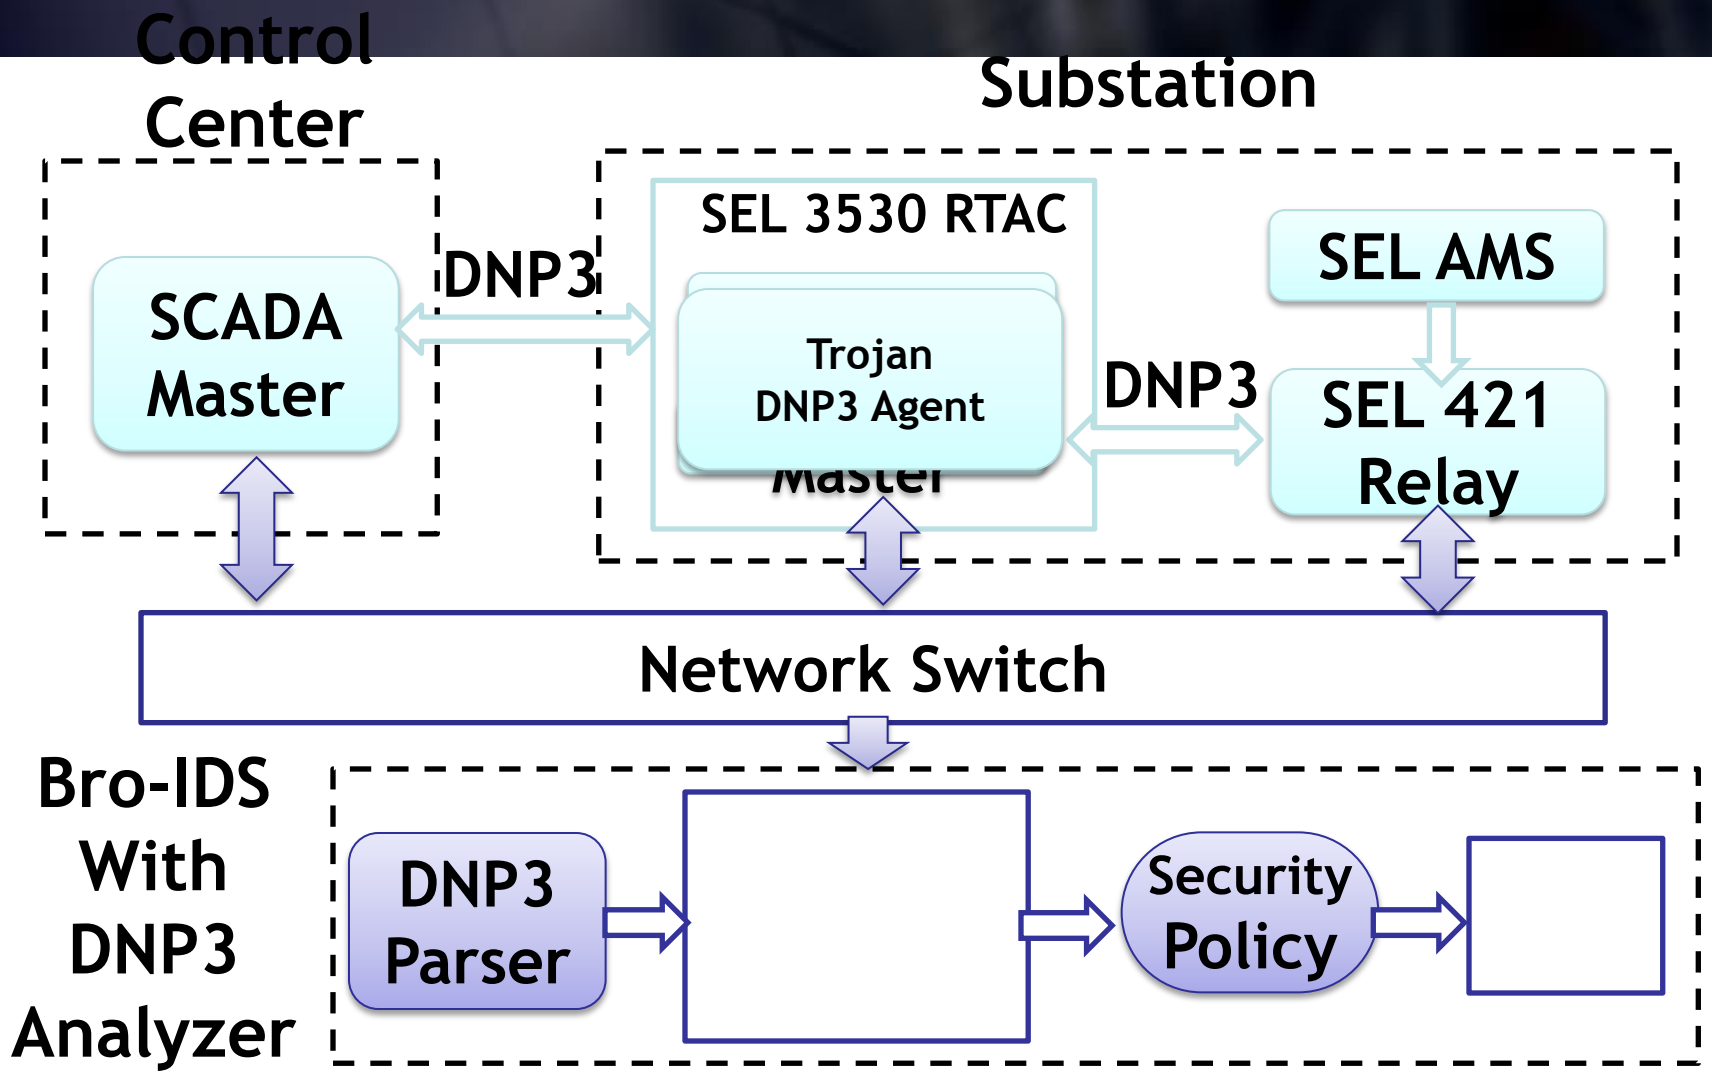

Semantic Analysis Framework

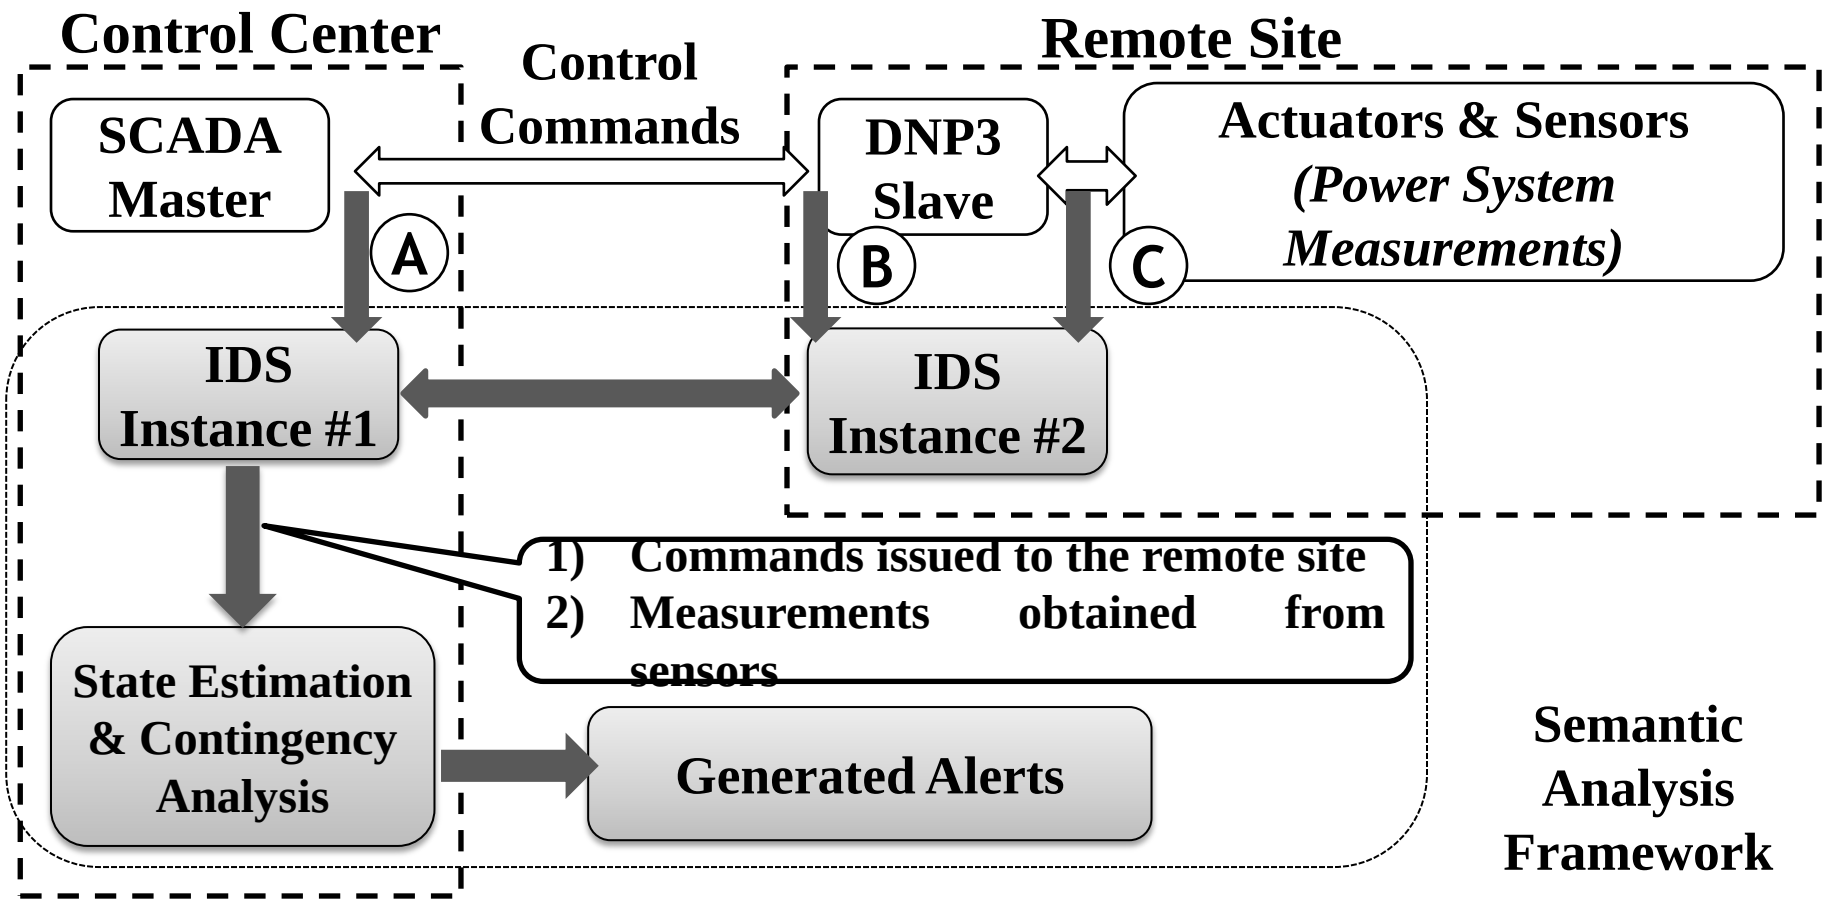

Issues: Quality of Security Monitoring

- **How does a monitor fail?**
  - Direct target of attacker
  - Missing invariants
  - Manipulated invariants

- **Robust monitoring**
  - Isolated from attackers
  - Robust invariants
  - Redundancy in Monitored Views
  - Compare the Monitored *invariants*

Major Challenges

- Develop and enforce security and reliability policies?
- Continuous orthogonal monitoring and invariance checking against attacks and failures
- Assured virtual environments
- Estimate, Validate Benchmark