Making ~~Forensic~~ ~~Attack~~ Event/forensic Analysis as Simple as Possible and No Simpler

Sean Peisert, UC Davis

Given at Schloss Dagstuhl — July 22, 2008

Electronic Voting Machines

- Need to be able to count ballots
- Need to be able to determine if and how a machine failed.
- Cannot allow a voter to indicate to an auditor who they are (vote selling)
- Cannot allow an auditor to determine who a voter is (voter coersion)
- This leads to a direct conflict. So how do we balance this?
  - Add noise
  - Enforce regularity

Existing Technical Solutions and the Insider Problem

- Access Control
- Intrusion Detection
- Anomaly-Based Detection
- Misuse-Based Detection
- Signature-Based Detection

Optimistic Access Control

- Security and usability are in conflict.
- Ideally, a system should block all forbidden actions and permit all allowed actions. (This is not feasible.)
- Policies can be binary (block access) or flexible (perform this countermeasure).
- Policies can be static (always do this) or dynamic (uh oh—an intruder)
- Many possible countermeasures exist
  - log
  - checkpoint/replay
  - make a particular partition read-only
- Many possible dynamic approaches exist
  - Use an a standard IPS
  - Incorporate external factors

So we need to focus on non-binary (e.g., *post mortem* analysis).

What is forensic Analysis?

- forensic analysis is the process of answering the questions:
  - How did an event take place?
  - What was the nature of the event?
  - What were the effects of the event?
- forensic analysis applies to arbitrary events. This can include attacks, but is not limited to attacks (e.g., mistakes).
- forensic analysis is not intrusion detection.
  - The goal of intrusion detection is to determine whether an attack occurred.

Transparent Society

(abbreviated from David Brin's ideas)

- Anyone can know anything.
- There is no privacy.
- It's better if everyone knows everything than if a few people know everything.
- "Watching the watchers"
- R. Heinlein: "'privacy laws' only make the bugs smaller."

Audit trails are...

- Is it is not well understood what forensic data is necessary, and there is no general solution to find that data.
- Data is often redundant, missing, vague, or misleading.
- Forensic analysis is worthless with bad data.
- We're wasting time, drawing bad conclusions, and making bad decisions.
- We need better data.
- A systematic approach to forensic logging gives better data and better analysis.

Current State

- Decent tools, but what problem do they solve?
  - file & filesystem analysis (Coroner's Toolkit, Sleuth Kit, EnCase, FTK)
  - syslog, tcpwrappers, Windows event logs
  - BSM
  - process accounting logs
  - IDS logs
  - packet sniffing

Forensics

- What do we need?
- What are we missing?

What are the assumptions for using current forensic tools?

- Often that there's only one person who had access to the machine.
- Often that the owner of the machine was in complete control (as opposed to malware).
- Probably a lot of other assumptions that we have no clue about.

For forensics, we need to...

- go back to the beginning.
- understand what the purpose of the analysis is
- understand what data can answer that purpose, with X% accuracy, and under a set of Y assumptions
- log the data
- give tools and techniques to an analyst to analyze that data

Art & Science

- But computer science can only answer part of it.
- Forensic analysis is an art, but there *are* scientific components. What are they?
  - Determining what to log
  - Determining relevance of logged data
    - what is relevant?
    - what is not relevant?
    - under what circumstances something might be relevant?
  - Using the results to constrain and correlate data.
  - *This can be measured, systematized and automated.*

Logging

- Two options:
  - Log everything (e.g., all non-deterministic events), and capture upon replay
  - Log selectively
    - *Ad hoc*
    - Systematic (e.g., based on security policies)

A Systematic Approach is Better

- Given system *S*, that records data *D*, what intrusions I_D can we understand with the data we have?
- Given intrusions *I'*, what additional data $D_{I'}$ do we need to record to analyze those intrusions?
- Given an arbitrary system defined by certain specifications, what information must be logged to detect violations of those specifications?

Laocoön

- *Laocoön: A Model of Forensic Logging*
- Attack graphs of goals.
- Goals can be attacker goals or defender goals (i.e., "security policies")
- Pre-conditions & post-conditions of events to accomplish goals.
- Method of translating those conditions into logging requirements.
- Logs are in a standardized and parseable format.
- Logged data can be at arbitrary levels of granularity.

Goals

- Premise: compute resources are cheap, human time is expensive.
- Understand the scope of the possible data, analyses, and conclusions.
- Be able to define (or place bounds on) what necessary information is present and what is missing.
- Assuming all potentially relevant information is recorded (e.g., by extrospection of a virtual machine), be able to correlate and prune the information necessary for a human to analyze.

Attack Graphs

- Intruder goals can be enumerated.
- Vulnerabilities, attacks, and exploits cannot (or in many cases, we would patch them, or they would inhibit usability).
- Defender goals can also be enumerated. They are called security polices.

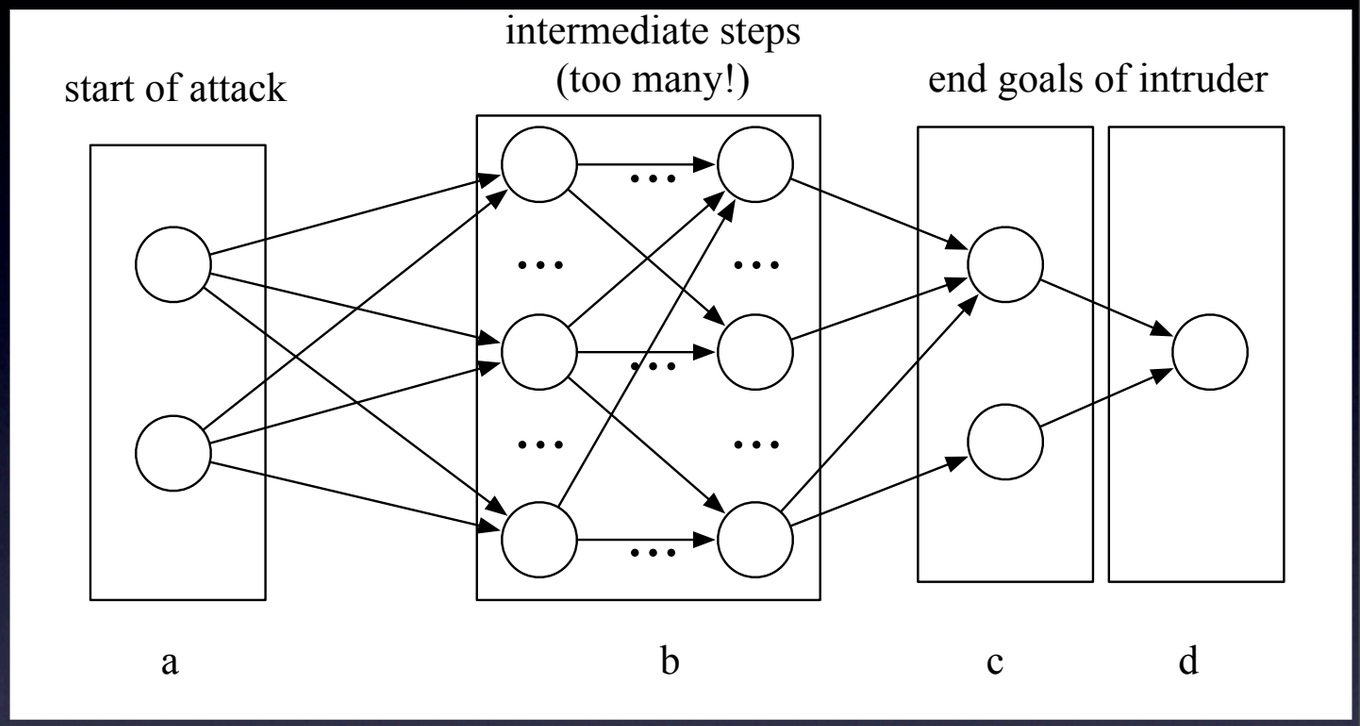

Security Policies

- Legal policies (HIPAA, Sarbanes-Oxley)
- Formal policies (Bell-LaPadula, Chinese Wall)
- Actual metrics
  - Severity (path length, time, difficulty)
  - Attack Surface Metric
  - Historically known vulnerabilities

Security Policies

- Security policies can be reverse-engineered or enforced, automatically.
- i.e., determine the current policy, and modify.
  - Policies can be binary (block access) or flexible (log something).
  - Policies can be static (always do this) or dynamic (uh oh—an intruder)
- Assumptions get in the way of security. What are they?

Applying Security Policies

- Applying Laocoön to security policies guides where to place instrumentation and what to log.
- The logged data needs to be correlated with a unique path identifier.
- Branches of a graph unrelated to the attack can be automatically pruned.
- Avoid recording data where events can be recreated because they are deterministic.

Pruning Paths

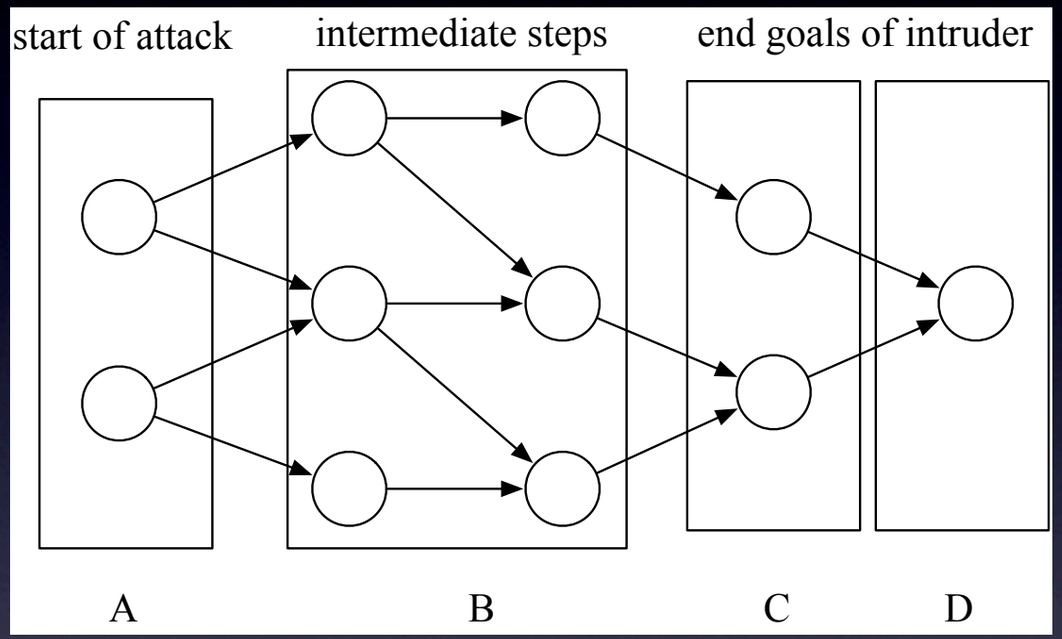

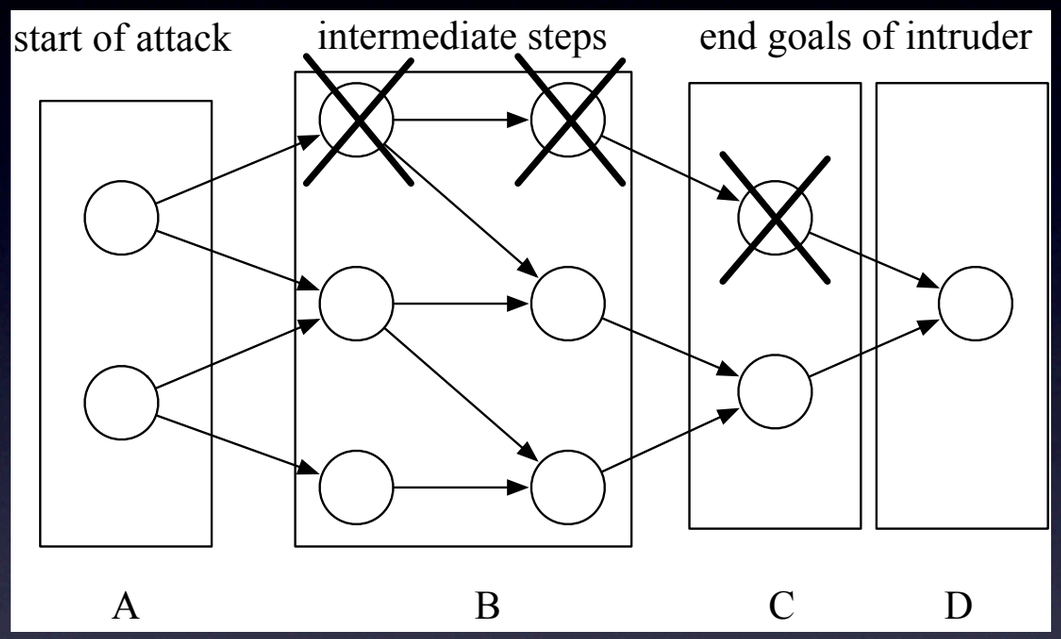

Pruning Paths

Aug 13 16:51:44 deimos kernel: Call 5: 9/0 = 495 syslogd(/etc/syslog.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 10/0 = 495 syslogd(/dev/console:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 11/0 = 495 syslogd(/var/log/messages:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 12/0 = 495 syslogd(/var/log/channel.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 13/0 = 495 syslogd(/var/log/security:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 14/0 = 495 syslogd(/var/log/auth.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 15/0 = 495 syslogd(/var/log/maillog:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 16/0 = 495 syslogd(/var/log/lpd-errs:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 17/0 = 495 syslogd(/var/log/xferlog:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 18/0 = 495 syslogd(/var/log/cron:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 19/0 = 495 syslogd(/var/log/debug.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 20/0 = 495 syslogd(/var/log/slip.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 deimos kernel: Call 5: 21/0 = 495 syslogd(/var/log/ppp.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 59: 0/0 = 932 ps((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 66: 932/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 0/2 = 932 ps(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1699750299:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 932 ps(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 932 ps(/lib/libm.so.4:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 932 ps(/lib/libkvm.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 932 ps(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 59: 0/0 = 933 grep((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 66: 933/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 932 ps(/dev/null:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 4/0 = 932 ps(/dev/null:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 5/0 = 932 ps(/etc/nsswitch.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 5/0 = 932 ps(/etc/spwd.db:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 6/0 = 932 ps(/etc/localtime:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 0/2 = 933 grep(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 572779291:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 933 grep(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 933 grep(/usr/lib/libgnuregex.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 933 grep(/usr/lib/libbz2.so.2:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 933 grep(/lib/libz.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 deimos kernel: Call 5: 3/0 = 933 grep(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 66: 934/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 59: 0/0 = 934 ls((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 0/2 = 934 ls(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1212716160:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 3/0 = 934 ls(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 3/0 = 934 ls(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 3/0 = 934 ls(/lib/libncurses.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 3/0 = 934 ls(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 3/0 = 934 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 4/0 = 934 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 deimos kernel: Call 5: 5/0 = 934 ls(..4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:50 deimos kernel: Call 5: 0/0 = 923 csh(/var/:4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 59: 0/0 = 935 ls((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 66: 935/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 0/2 = 935 ls(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -573042789:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 3/0 = 935 ls(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 3/0 = 935 ls(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 3/0 = 935 ls(/lib/libncurses.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 3/0 = 935 ls(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 3/0 = 935 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 deimos kernel: Call 5: 4/0 = 935 ls(/var/log/:4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 deimos kernel: Call 59: 0/2 = 936 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 deimos kernel: Call 59: 0/0 = 936 du((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 deimos kernel: Call 5: 0/2 = 936 du(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 deimos kernel: Call 5: 3/0 = 936 du(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 deimos kernel: Call 5: 3/0 = 936 du(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 deimos kernel: Call 5: 3/0 = 936 du(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)

Pruning Paths

Aug 13 16:51:44 daimos kernel: Call 5: 9/0 = 495 syslogd(/etc/syslog.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 10/0 = 495 syslogd(/dev/console:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 11/0 = 495 syslogd(/var/log/messages:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 12/0 = 495 syslogd(/var/log/channel.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 13/0 = 495 syslogd(/var/log/security:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 14/0 = 495 syslogd(/var/log/auth.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 15/0 = 495 syslogd(/var/log/maillog:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 16/0 = 495 syslogd(/var/log/lpd-errs:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 17/0 = 495 syslogd(/var/log/xferlog:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 18/0 = 495 syslogd(/var/log/cron:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 19/0 = 495 syslogd(/var/log/debug.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 20/0 = 495 syslogd(/var/log/slip.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 21/0 = 495 syslogd(/var/log/ppp.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 59: 0/0 = 932 ps((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 66: 932/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 0/2 = 932 ps(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1699750299:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/lib/libm.so.4:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/lib/libkvm.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 59: 0/0 = 933 grep((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 66: 933/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/dev/null:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 4/0 = 932 ps(/dev/null:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 5/0 = 932 ps(/etc/nsswitch.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 5/0 = 932 ps(/etc/spwd.db:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 6/0 = 932 ps(/etc/localtime:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 0/2 = 933 grep(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 572779291:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/usr/lib/libgnuregex.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/usr/lib/libbz2.so.2:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/lib/libz.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 66: 934/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 59: 0/0 = 934 ls((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 0/2 = 934 ls(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1212716160:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/lib/libncurses.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 4/0 = 934 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 5/0 = 934 ls(..4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:50 daimos kernel: Call 5: 0/0 = 923 csh(/var/:4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 59: 0/0 = 935 ls((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 66: 935/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 0/2 = 935 ls(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -573042789:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/lib/libncurses.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 4/0 = 935 ls(/var/log/:4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 59: 0/2 = 936 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 59: 0/0 = 936 du((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 0/2 = 936 du(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 3/0 = 936 du(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 3/0 = 936 du(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 3/0 = 936 du(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 66: 936/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)

Graph 62

Node 1

Node 3

Node 4

Node 5

Pruning Paths

Aug 13 16:51:44 daimos kernel: Call 5: 9/0 = 495 syslogd(/etc/syslog.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 10/0 = 495 syslogd(/dev/console:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 11/0 = 495 syslogd(/var/log/messages:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 12/0 = 495 syslogd(/var/log/channel.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 13/0 = 495 syslogd(/var/log/security:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 14/0 = 495 syslogd(/var/log/auth.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 15/0 = 495 syslogd(/var/log/maillog:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 16/0 = 495 syslogd(/var/log/lpd-errs:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 17/0 = 495 syslogd(/var/log/xferlog:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 18/0 = 495 syslogd(/var/log/cron:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 19/0 = 495 syslogd(/var/log/debug.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 20/0 = 495 syslogd(/var/log/slip.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:44 daimos kernel: Call 5: 21/0 = 495 syslogd(/var/log/ppp.log:9:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 1, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 59: 0/0 = 932 ps((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 66: 932/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 0/2 = 932 ps(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1699750299:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/lib/libm.so.4:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 932 ps(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 66: 933/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 4/0 = 932 ps(/dev/null:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 5/0 = 932 ps(/etc/nsswitch.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 6/0 = 932 ps(/etc/localtime:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 0/2 = 933 grep(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 572779291:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/usr/lib/libgnuregex.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:46 daimos kernel: Call 5: 3/0 = 933 grep(/lib/libz.so.3:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 66: 934/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 59: 0/0 = 934 ls((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 0/2 = 934 ls(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1212716160:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/lib/libncurses.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 3/0 = 934 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 4/0 = 934 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:47 daimos kernel: Call 5: 5/0 = 934 ls(..4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:50 daimos kernel: Call 5: 0/0 = 923 csh(/var/:4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 66: 935/0 = 923 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 922, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 0/2 = 935 ls(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -573052789:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/lib/libncurses.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 3/0 = 935 ls(..0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 1001:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:51 daimos kernel: Call 5: 4/0 = 935 ls(/var/log/:4:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 59: 0/2 = 936 csh((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 59: 0/0 = 936 du((null):0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 0/2 = 936 du(/etc/libmap.conf:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: -1:2, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 3/0 = 936 du(/var/run/ld-elf.so.hints:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 3/0 = 936 du(/lib/libutil.so.5:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)
Aug 13 16:51:54 daimos kernel: Call 5: 3/0 = 936 du(/lib/libc.so.6:0:(null):0) RUID/E/S: 0/0/0, RGID/S: 0/0, FilUID: 0:0, PPID: 923, UPI(0,0,0,0)

Graph 62

Node 1

Exploit

Node 3

Node 4

Node 5

Complex Attack Graph

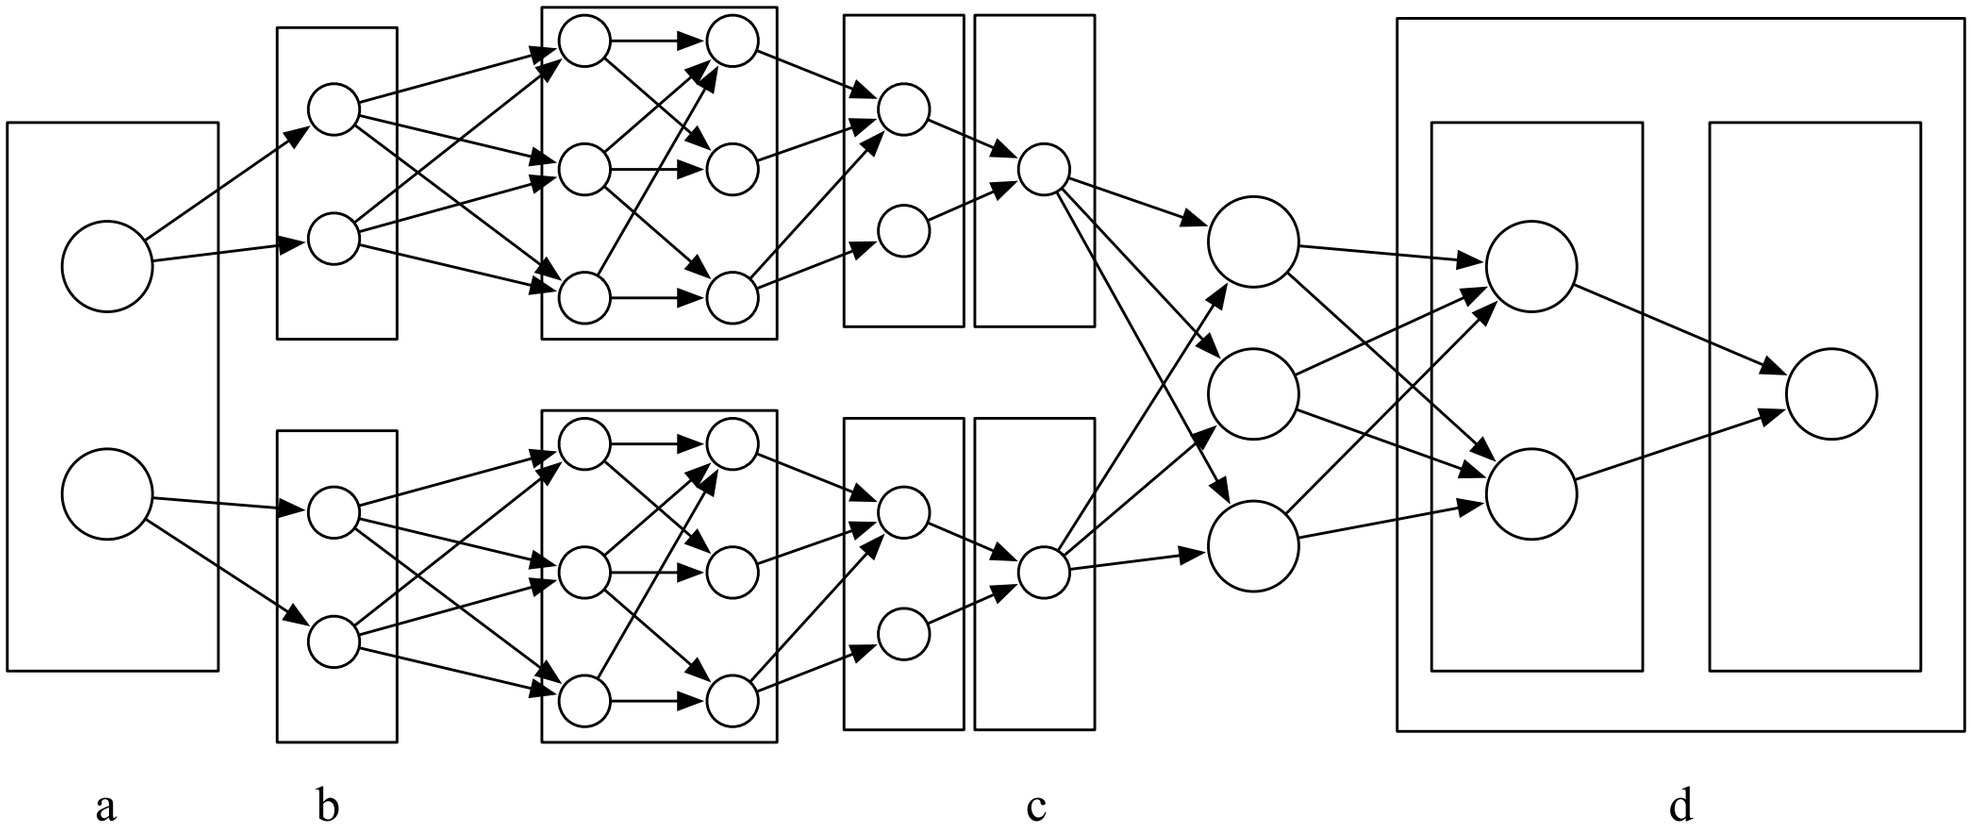

Summary

- Forensics, attack analysis, logging, and auditing are broken.
- We have developed methods to correlate and constrain data that needs to be analyzed.
- We have developed methods for logging based on known vulnerabilities.
- We have developed methods for integrating societal needs (e.g., law) with forensic logging and auditing capabilities.

Thank you

- Questions?
- Sean Peisert
  - peisert@cs.ucdavis.edu
  - http://www.sdsc.edu/~peisert/