Lecture 20: Audit

CS 181S

November 26, 2018

Classes of Countermeasures

- **Authentication:** mechanisms that bind principals to actions

- **Authorization:** mechanisms that govern whether actions are permitted

- **Audit:** mechanisms that record and review actions

Uses of audit

- **Deterrence through accountability:** deter misbehavior

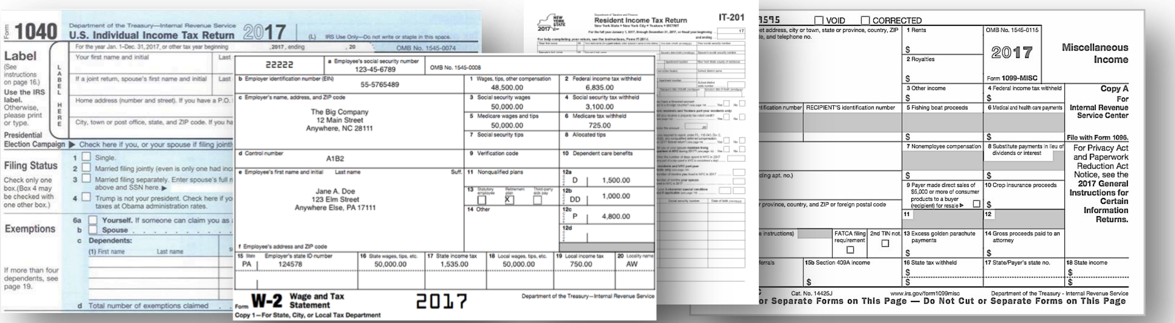

1040
U.S. Individual Income Tax Return 2017
Resident Income Tax Return
IT-201
VOID
CORRECTED
2017
Form 1099-MISC
Miscellaneous Income
Copy A
For Internal Revenue Service Center

22222	a Employee's social security number 123-45-6789	OMB No. 1545-0008
b Employer identification number (EIN) 55-5765489	1 Wages, tips, other compensation 48,500.00	2 Federal income tax withheld 6,835.00
c Employer's name, address, and ZIP code The Big Company 12 Main Street Anywhere, NC 28111	3 Social security wages 50,000.00	4 Social security tax withheld 3,100.00
	5 Medicare wages and tips 50,000.00	6 Medicare tax withheld 725.00
	7 Social security tips	8 Allocated tips
d Control number A1B2	9 Verification code	10 Dependent care benefits
e Employee's first name and initial Last name Suff. Jane A. Doe 123 Elm Street Anywhere Else, PA 17111	11 Nonqualified plans	12a D 1,500.00
	13 Statutory employee □ Retirement plan ☒ Third-party sick pay □	12b DD 1,000.00
	14 Other	12c P 4,800.00
f Employee's address and ZIP code		12d

15 State	Employer's state ID number	16 State wages, tips, etc.	17 State income tax	18 Local wages, tips, etc.	19 Local income tax	20 Locality name
PA	124578	50,000.00	1,535.00	50,000.00	750.00	AW

Form **W-2** Wage and Tax Statement 2017 Department of the Treasury—Internal Revenue Service

Copy 1—For State, City, or Local Tax Department

- **Detection and recovery:** determine what happened and how to recover

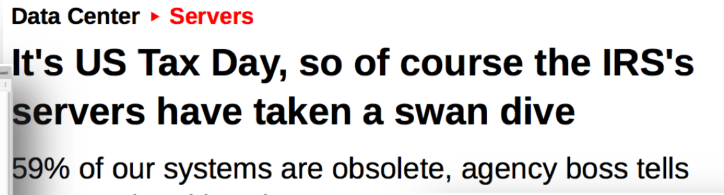

Data Center ▸ Servers

It's US Tax Day, so of course the IRS's servers have taken a swan dive

59% of our systems are obsolete, agency boss tells congressional hearing

By Thomas Claburn in San Francisco 17 Apr 20

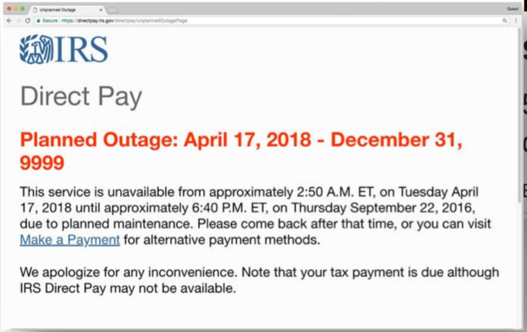

IRS

Direct Pay

Planned Outage: April 17, 2018 - December 31, 9999

This service is unavailable from approximately 2:50 A.M. ET, on Tuesday April 17, 2018 until approximately 6:40 P.M. ET, on Thursday September 22, 2016, due to planned maintenance. Please come back after that time, or you can visit Make a Payment for alternative payment methods.

We apologize for any inconvenience. Note that your tax payment is due although IRS Direct Pay may not be available.

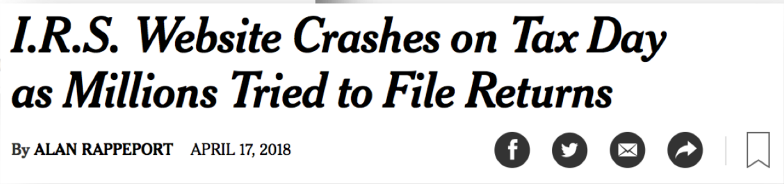

I.R.S. Website Crashes on Tax Day as Millions Tried to File Returns

By ALAN RAPPEPORT APRIL 17, 2018

- **Problem monitoring:** real-time intelligence

Audit tasks

- **Recording:**
  - what to log
  - what not to log
  - how to log
    - locally
    - remotely
  - how to protect the log
- **Reviewing:**
  - automated analysis
  - manual exploration

WHAT TO LOG

What to log?

Example: US State Department pilot program (1980s)

- Requirements:
  - log every transaction related to protected electronic documents
  - system administrator reviews log daily to search for malicious behavior
- Experiment:
  - test system for 5 users, 10 minutes
  - audit log was a stack of paper
  - real system would have been 1000s of users working 24/7
- Lessons learned:
  - logging and review of everything by a human is impractical
  - need to reduce information logged: log reduction
  - need automated review

States vs. events

- **States:** data, *what the system is*
  - backup, or more
  - survive power failures, crashes, attacks
  - what state? memory, disk, network, ...
  - consistent snapshot of distributed system is hard

- **Events:** actions, *how the system came to be*
  - login, access to protected resource, elevation and attenuation of privileges, ...
  - our focus
  - which events?

Recall: Security requirements

- **Functional requirement:** something system should do
  - e.g., allow people to cash checks
- **Security goal:** something system should/shouldn't do
  - e.g., prevent loss of revenue through bad checks
- **Security requirement:** constraint on functional requirement to achieve goal
  - e.g., check must be drawn on bank where being cashed, or person cashing must be customer at that bank and deposit in their account

Events to log

- **Any event that involves a security requirement**
  - Fact that requirement was checked
  - Whether it was met or not
  - The information that led to that decision
- Typically involves the gold standard...
  - whether a **principal was authenticated,** or
  - whether an **action was authorized**

Orange Book logging

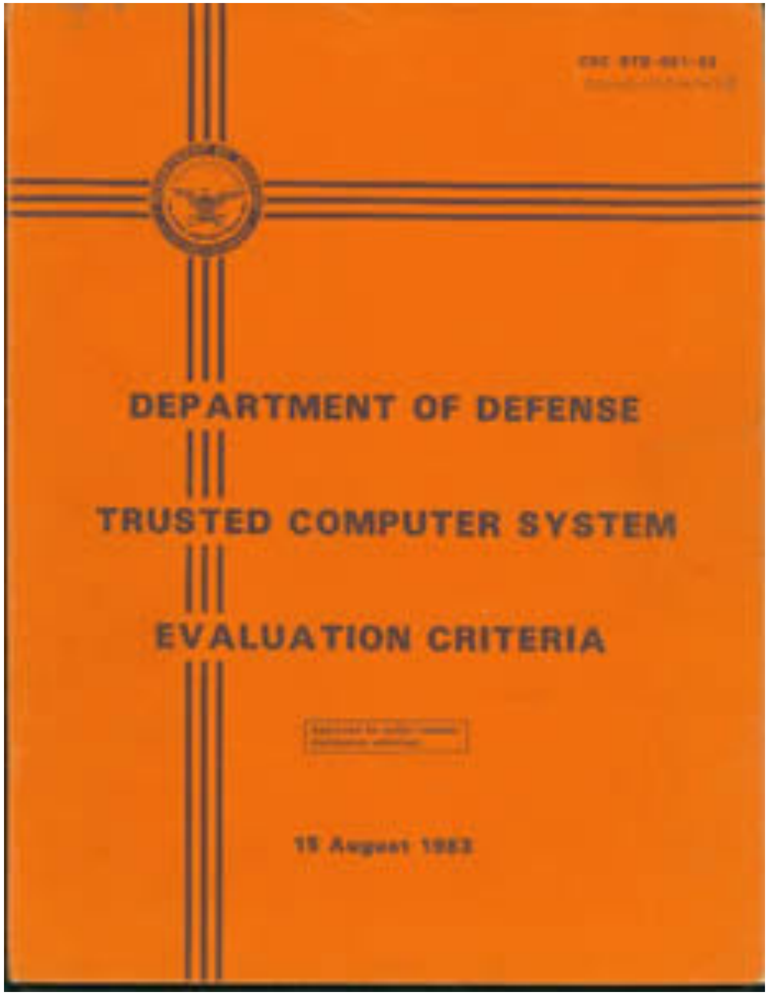

For minimal C2 level certification:

- **Events** to log:
  - Use of identification and authentication mechanisms
  - Introduction of objects into a user's address space (e.g., file open, program initiation)
  - Deletion of objects
  - Actions taken by computer operators and system administrators and/or system security officers

Orange Book logging

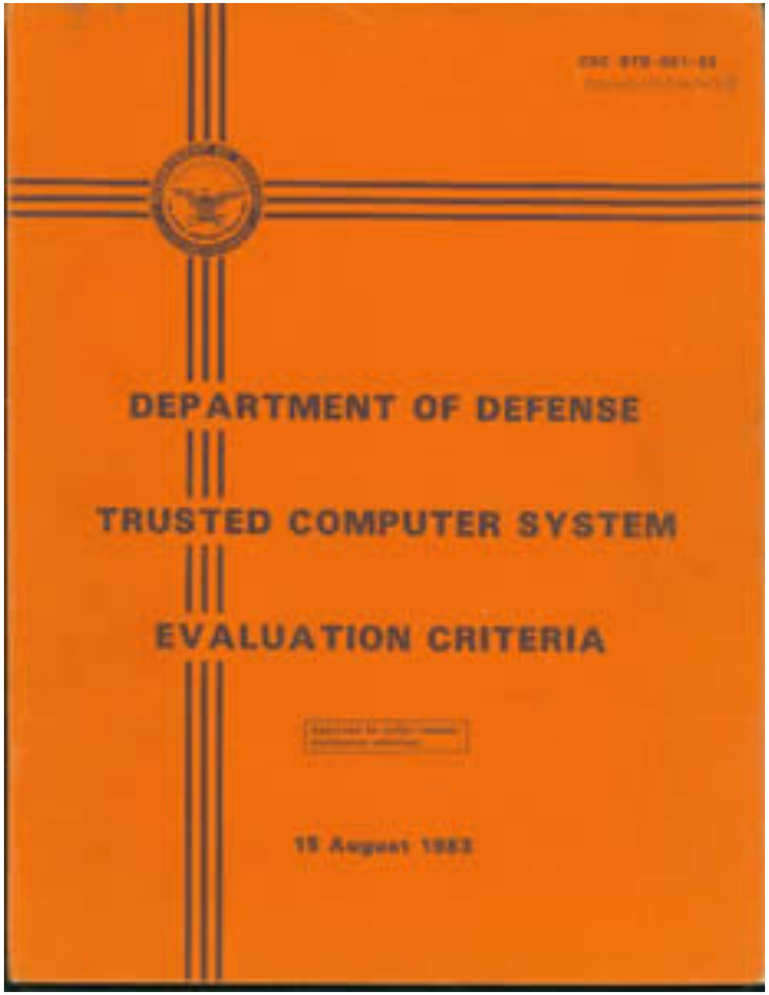

DEPARTMENT OF DEFENSE

TRUSTED COMPUTER SYSTEM

EVALUATION CRITERIA

15 August 1983

For minimal C2 level certification:

- **What** to log:
  - Date and time of the event
  - User
  - Type of event
  - Success or failure of the event
  - For identification/authentication events: origin of request
  - For events involving objects: name of the object

What not to log

- Some information might be too sensitive for log files:
  - plaintext keys, passwords
  - the details of company's shiny new product
  - the GPS coordinates of undercover secret agents

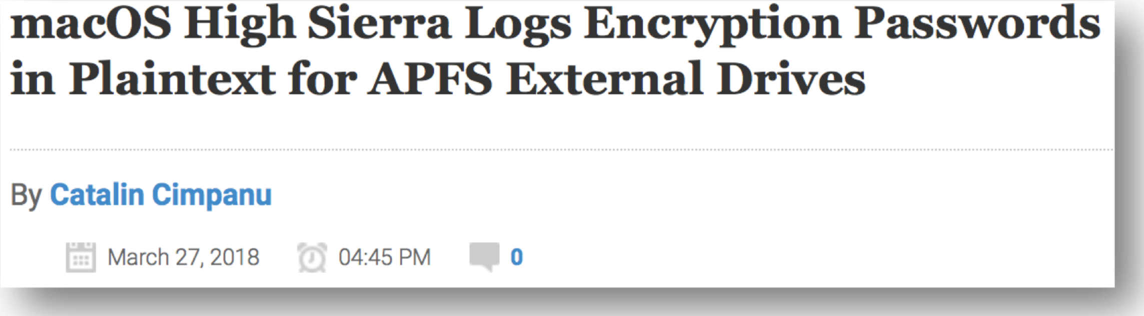

macOS High Sierra Logs Encryption Passwords in Plaintext for APFS External Drives

By Catalin Cimpanu

March 27, 2018 04:45 PM 0

- Possibilities:
  - log it anyway, protect the log
  - sanitize log

Sanitization

Protect confidential information in log

- by **deleting**
- by **modifying**
  - e.g., replace with user names with pseudonyms, keep separate protected map between names and pseudonyms

Sanitization

- Before writing to log:
  - **Pro:** protects users from system administrators; maybe surveillance warranted only with probable cause
  - **Con:** have to decide in advance, as part of system design, what information to keep vs. discard
- After writing to log:
  - **Con:** confidentiality of log must be (more) protected
  - **Pro:** can decide afterwards what information to discard, perhaps even redact logs and send to 3rd party for analysis

Examples: CMS and Sakai

Example: CMS

Details logged:

- Event type
- Acting NetID
- Acting IP address
- Affected NetIDs
- Simulated NetID
- Assignment, if any
- Event details (no sanitization of grades)

HOW TO LOG

Say what you mean

Main principle: Every log entry should say what it means

- Interpretation of log entry should depend only upon content of log entry
- Hence reviewer can recover meaning without needing to assume or supply any context

Log file format

- Keeping log files in standard format enables...
  - Reuse of tools for log analysis
  - Correlation across logs from multiple applications
- Standard formats:
  - Common Log Format (used by web servers)
  - syslog (used by Unix)
    - originated with sendmail
    - became a *de facto* standard
    - then standardized by IETF: RFC 5424
    - examples: take a look in your local /var/log directory

Common Log Format

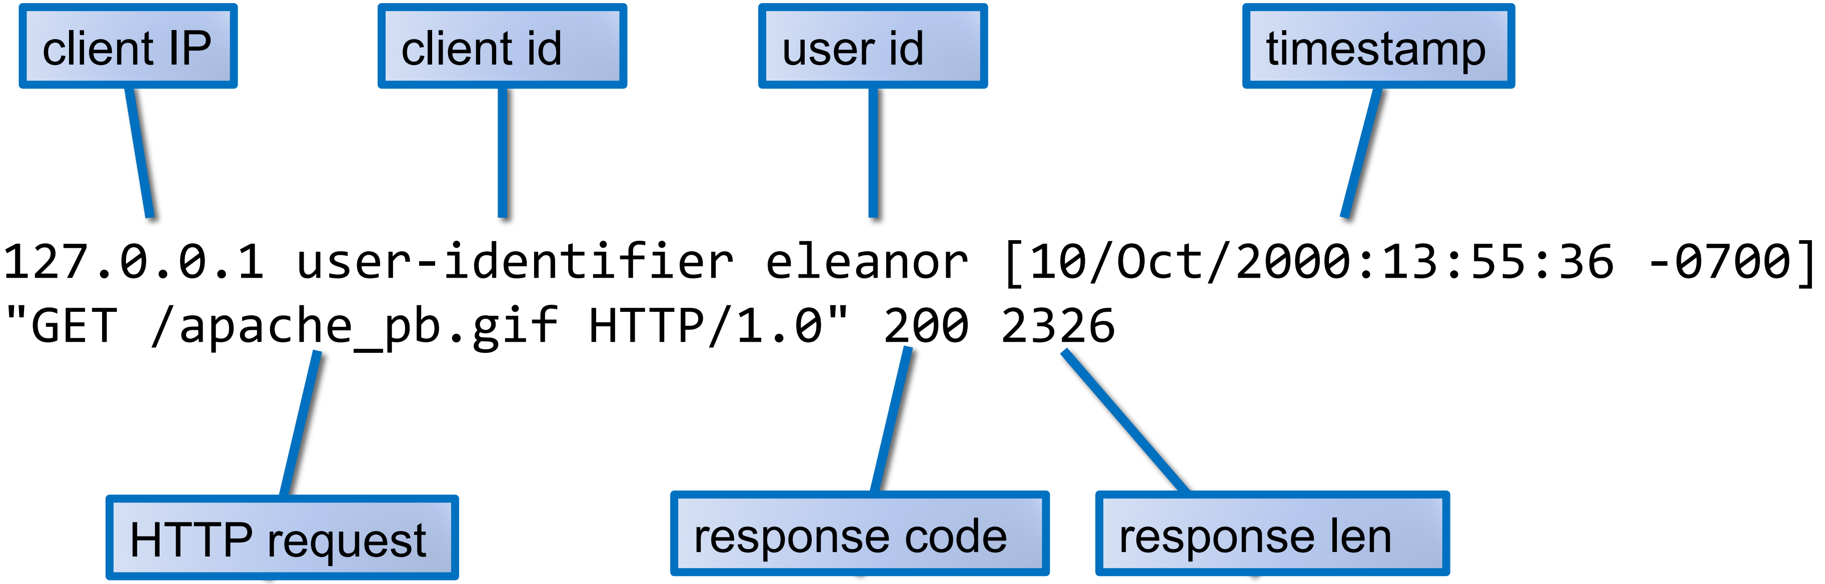

client IP

client id

user id

timestamp

```
127.0.0.1 user-identifier eleanor [10/Oct/2000:13:55:36 -0700]
"GET /apache_pb.gif HTTP/1.0" 200 2326
```

HTTP request

response code

response len

syslog example message

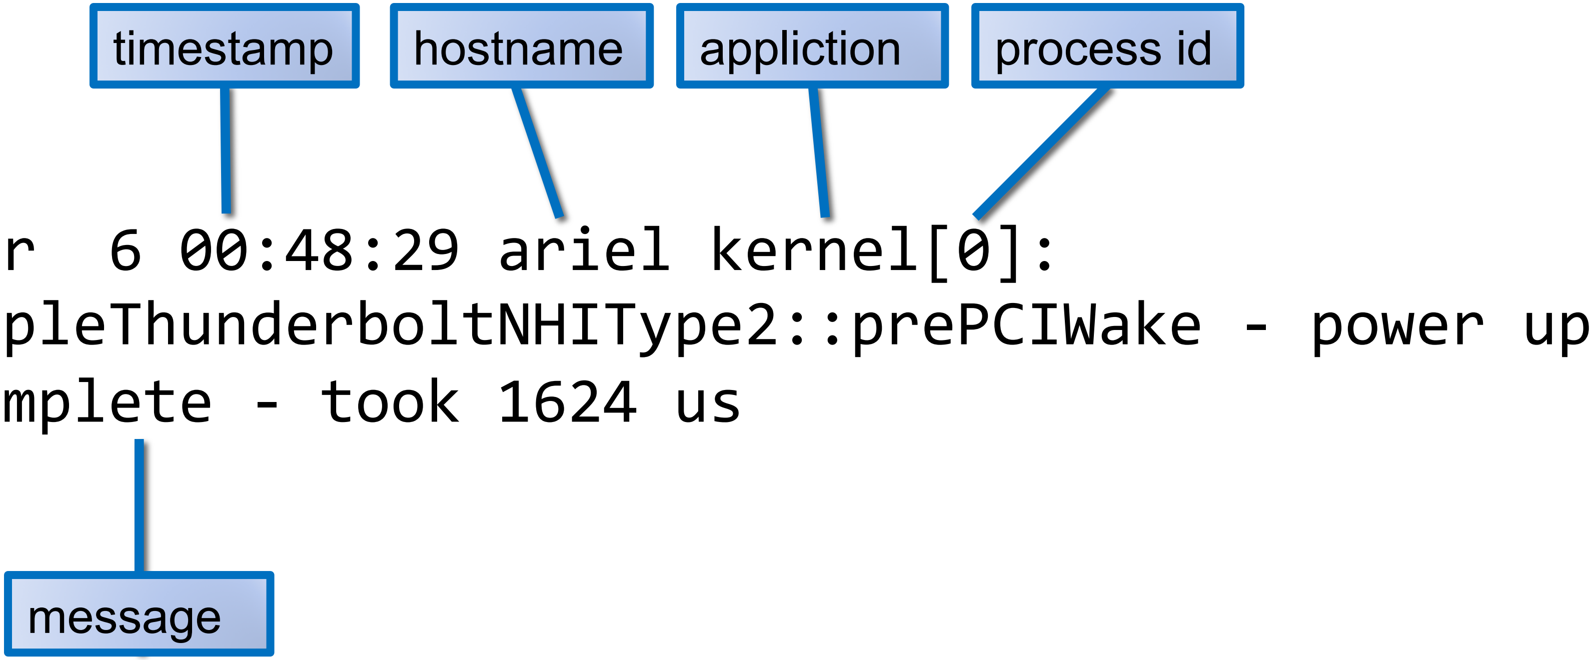

timestamp

hostname

appliction

process id

```
Mar  6 00:48:29 ariel kernel[0]:
AppleThunderboltNHIType2::prePCIWake - power up
complete - took 1624 us
```

message

Log space

What happens if log size grows too large?

- Halt system
- Overwrite previous entries
- Stop logging

SECURING THE LOG

Approaches to Securing Audit Log

- Limit access to log files
- Transmit entries to remote audit server
- Use cryptography

Limit Access to Log files

- least privilege
- limit who can read
- limit how principals can write (append-only for most users)

Remote Audit Servers

- how often?
- how secure log entries en route?

syslog architecture

- Originators: source of messages
  - might duplicate to multiple relays
- Relays: forward messages
  - might filter or duplicate messages
- Collectors: sink of messages
  - might collect from many sources

syslog architecture

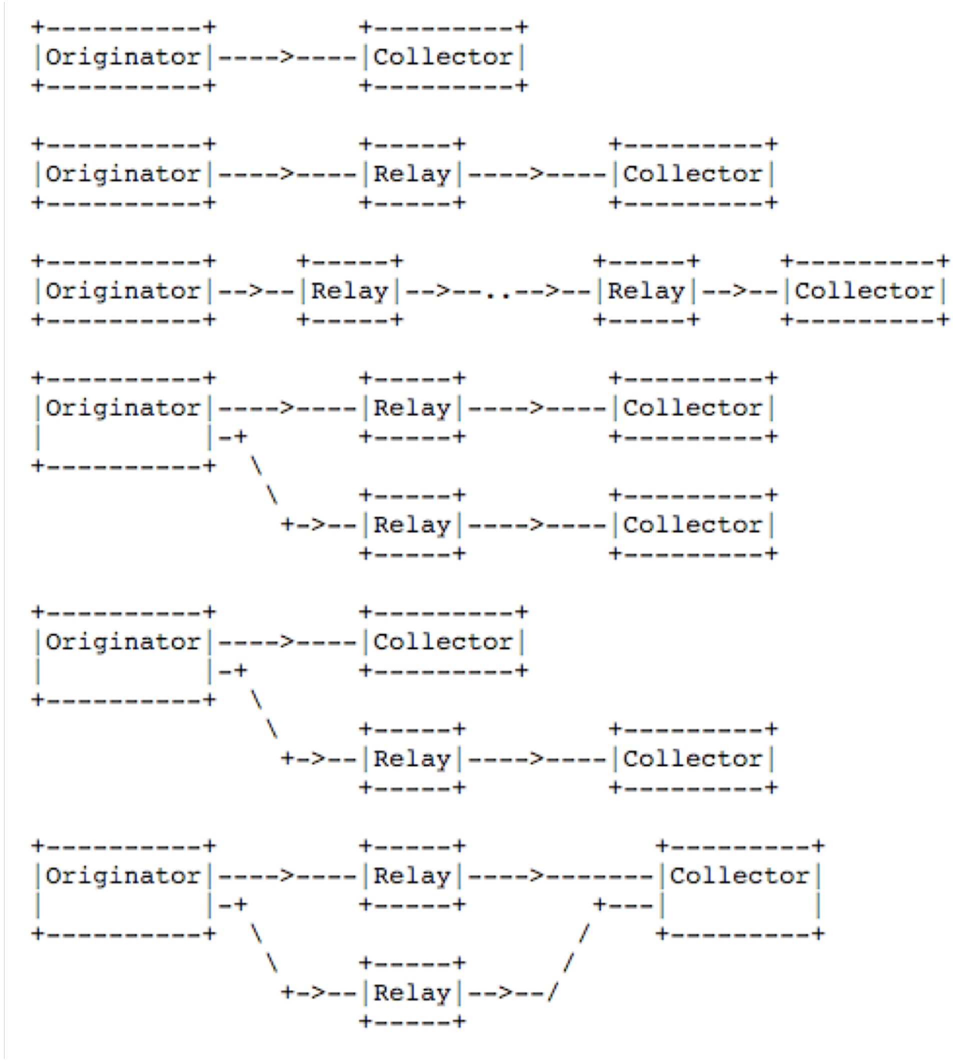

```
+----------+         +---------+
|Originator|---->----|Collector|
+----------+         +---------+

+----------+         +-----+         +---------+
|Originator|---->----|Relay|---->----|Collector|
+----------+         +-----+         +---------+

+----------+     +-----+            +-----+     +---------+
|Originator|-->--|Relay|-->--..-->--|Relay|-->--|Collector|
+----------+     +-----+            +-----+     +---------+

+----------+         +-----+         +---------+
|Originator|---->----|Relay|---->----|Collector|
|          |-+       +-----+         +---------+
+----------+  \
               \     +-----+         +---------+
                +->--|Relay|---->----|Collector|
                     +-----+         +---------+

+----------+         +---------+
|Originator|---->----|Collector|
|          |-+       +---------+
+----------+  \
               \     +-----+         +---------+
                +->--|Relay|---->----|Collector|
                     +-----+         +---------+

+----------+         +-----+            +---------+
|Originator|---->----|Relay|---->-------|Collector|
|          |-+       +-----+        +---|         |
+----------+  \                    /    +---------+
               \     +-----+      /
                +->--|Relay|-->--/
                     +-----+
```

Security concerns with syslog

Base syslog protocol has no security goals

- Recommended to use SSL to protect communication channel
- Nonetheless, receivers are permitted to truncate or drop messages
- Even with SSL, end-to-end integrity of messages from originator to collector not guaranteed
  - Concerns include provenance, message integrity, replays, sequencing, detection of missing messages
  - Digital signatures provide solution [RFC 5848]

Securing the log with crypto

- **Threat:** Attacker who compromises host that stores log. Attacker can read/write log file and can access secret keys
- **Harm:** log can be read, modified, deleted
- **Vulnerability:** log protected only by access control mechanisms on host (prior to archiving on remote server)

Securing the log with crypto

- **System:**
  - machine M maintains a local log
  - periodically M synchs log to trusted remote log server S
  - might be very long periods between synch: if short periods are possible, no real need for this protocol
- **Goals:** assume attacker compromises M at time t...
  - Contents of log messages entered before t are not disclosed to anyone who can read log at M (Confidentiality)
  - Contents of log messages and their sequence before time t cannot be changed in a way that is undetectable by S (Integrity)
- **Countermeasure:** cryptography: use iterated hashing: H(H(H(...H(v)...))) to create tamper-resistant log

Audit tasks

- **Recording:**
  - what to log
  - what not to log
  - how to protect the log
- **Reviewing:**
  - manual exploration
  - automated analysis

REVIEWING THE LOG

Manual review

- Enable administrators to explore logs and look for {states, events}
- **Issues:**
  - Designers might not have anticipated the right {states, events} to record
  - Visualization, query, expressivity (HCI/DB issues)
  - Correlation amongst multiple logs

Interfaces

- **Flat text** [example: last time's syslog]
- **Hypertext**
- **DBMS** [example: queries in CMS]
- **Visualization tools**

Techniques

- Temporal replay: animate what happened when
- Slice: display minimal set of log events that affect a given object

Automated review and response

- **Review:** detect suspicious behavior that looks like an attack, or detect violations of explicit policy
  - Custom-built systems
  - Classic AI techniques like training neural nets, expert systems, etc.
  - Modern applications of machine learning
- **Response:** report, take action