

# Logging

Portions courtesy Ellen Liu

# Outline

- Introduction
- Finding log files
- Syslog: the system event logger
- Linux “logrotate” tool
- Condensing log files to useful information
- Logging policies

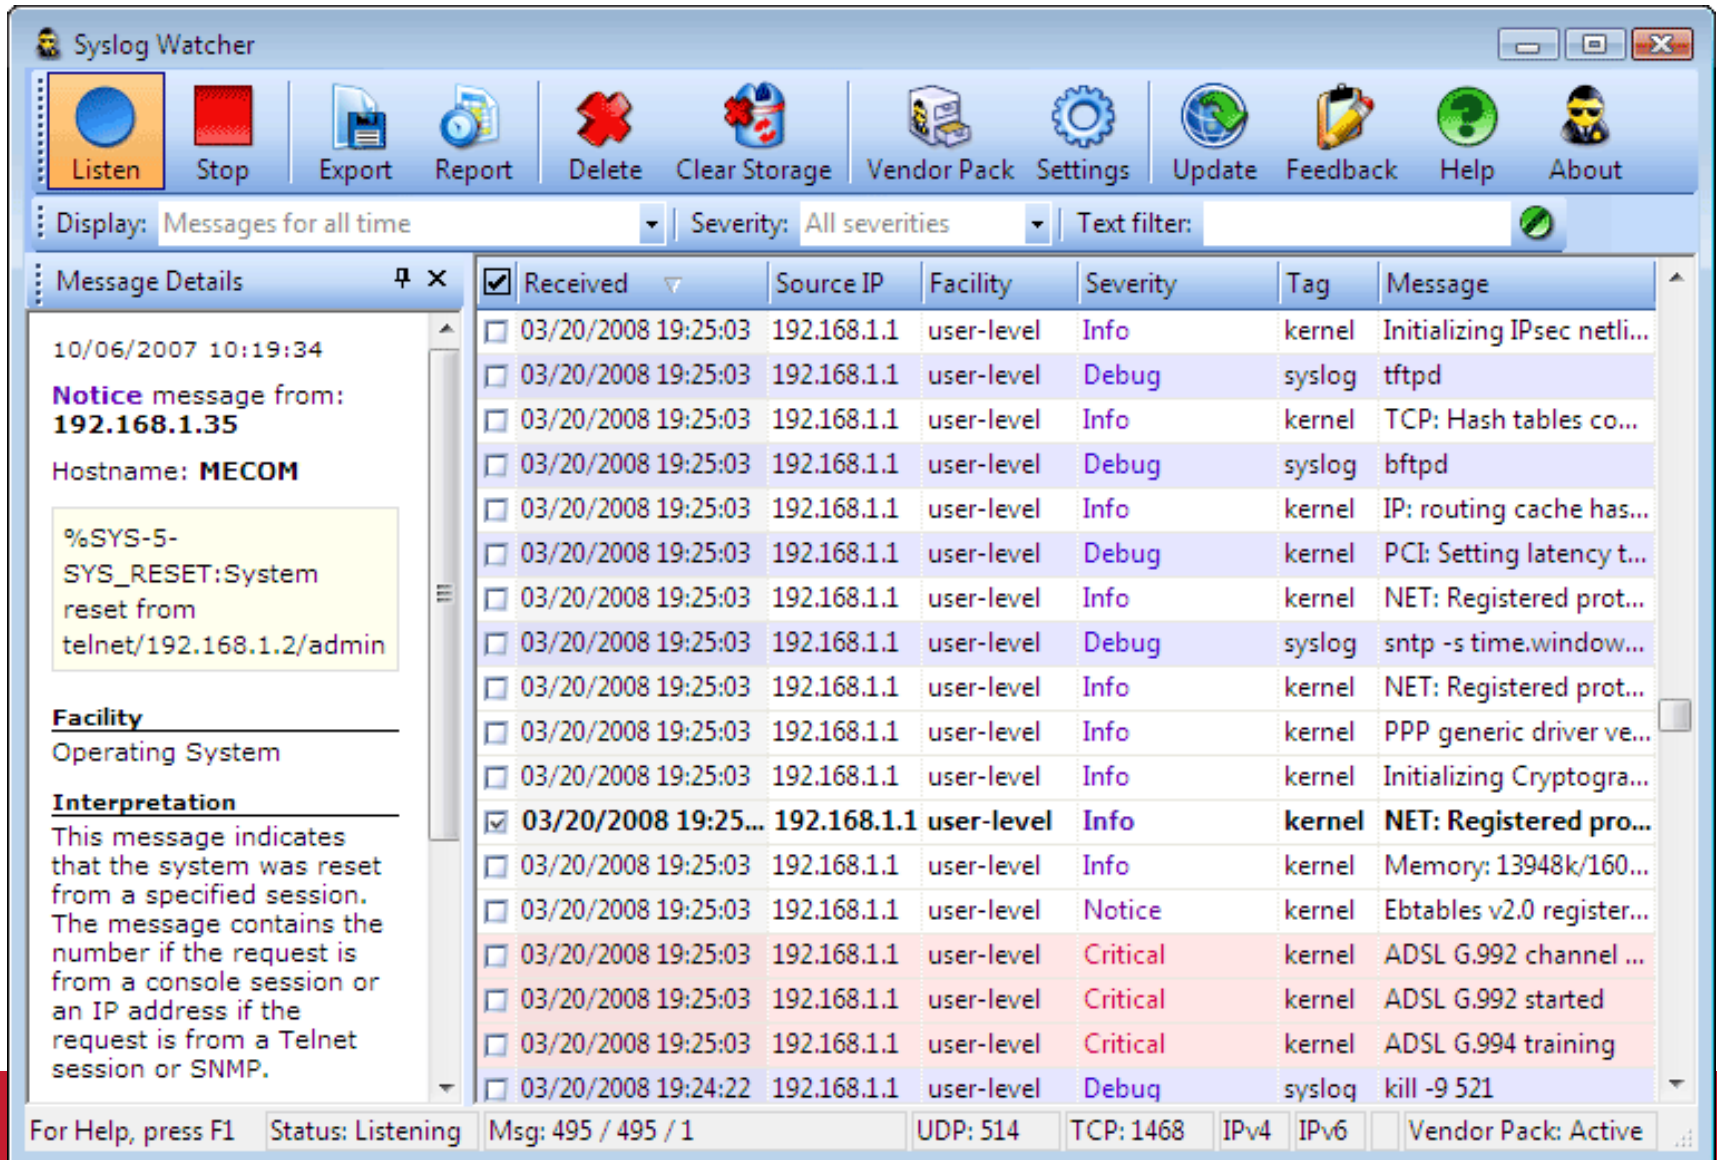
# Who and Why

- System daemons, kernel, and utilities produce log data onto disks
- Most log data has limited useful life, needs to be summarized, compressed, archived, then removed
- Access and audit data are needed by government regulations and company policies
- Logs also reveal configuration problems

# Logs

- A log event is captured as a single line of text
  - time and date, type and severity of the event, etc., often separated by spaces, tabs, or punctuation
- Logs are plaintext files, can be processed by shell commands and shell scripts
- There are also log management tools that rotate, compress, and monitor log files daily or weekly

# syslog messages



**Syslog Watcher**

Listen Stop Export Report Delete Clear Storage Vendor Pack Settings Update Feedback Help About

Display: Messages for all time Severity: All severities Text filter:

**Message Details**

10/06/2007 10:19:34

**Notice** message from:  
**192.168.1.35**

Hostname: **MECOM**

%SYS-5-  
SYS\_RESET:System  
reset from  
telnet/192.168.1.2/admin

**Facility**  
Operating System

**Interpretation**  
This message indicates that the system was reset from a specified session. The message contains the number if the request is from a console session or an IP address if the request is from a Telnet session or SNMP.

| <input checked="" type="checkbox"/> | Received             | Source IP   | Facility   | Severity | Tag    | Message                     |
|-------------------------------------|----------------------|-------------|------------|----------|--------|-----------------------------|
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | Initializing IPsec netli... |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Debug    | syslog | tftpd                       |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | TCP: Hash tables co...      |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Debug    | syslog | bftpd                       |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | IP: routing cache has...    |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Debug    | kernel | PCI: Setting latency t...   |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | NET: Registered prot...     |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Debug    | syslog | sntp -s time.window...      |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | NET: Registered prot...     |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | PPP generic driver ve...    |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | Initializing Cryptogra...   |
| <input checked="" type="checkbox"/> | 03/20/2008 19:25:... | 192.168.1.1 | user-level | Info     | kernel | NET: Registered pro...      |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Info     | kernel | Memory: 13948k/160...       |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Notice   | kernel | Ebttables v2.0 register...  |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Critical | kernel | ADSL G.992 channel ...      |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Critical | kernel | ADSL G.992 started          |
| <input type="checkbox"/>            | 03/20/2008 19:25:03  | 192.168.1.1 | user-level | Critical | kernel | ADSL G.994 training         |
| <input type="checkbox"/>            | 03/20/2008 19:24:22  | 192.168.1.1 | user-level | Debug    | syslog | kill -9 521                 |

For Help, press F1 Status: Listening Msg: 495 / 495 / 1 UDP: 514 TCP: 1468 IPv4 IPv6 Vendor Pack: Active

# IT Standards & Industry Regulations

- COBIT
  - A set of best practices framework for information technology (IT) management
- ISO 27002
  - Provides best practice recommendations on information security management
- Require sites to maintain a centralized, hardened, enterprise-wide repository for logs, with NTP time stamps and a strict retention schedule

# Finding Log Files

- **Names:** maillog, ftp.log, lpd-errs, console\_log, ...
- For Linux, by default most are found in `/var/log`, `/var/adm`
- Some common log files:

| File     | Program    | Where | Freq | Contents                    |
|----------|------------|-------|------|-----------------------------|
| acpid    | acpid      | F     | 64K  | Power-related events        |
| boot.log | rc scripts | F     | M    | Output of startup scripts   |
| cron     | cron       | S     | W    | cron executions and errors  |
| faillog  | login      | H     | W    | Unsuccessful login attempts |
| httpd/*  | httpd      | F     | D    | Apache logs                 |
| yum.log  | yum        | F     | M    | Package management log      |

Where (filename source) - S: syslog, H: hardwired, F: configuration file  
Freq (freq. of cleanup) - D: Daily, W: Weekly, M: Monthly, Size-based

# Log Permissions and Syslog

- Log files are normally owned by root
  - Occasionally by less privileged httpd, mysqld, etc.
- Sensitive logs need tight permissions. Others can be set to world-readable
- **Syslog**: an integrated system to concentrate logs
  - On UNIX/Linux systems
  - syslogd daemon
  - configuration file: </etc/syslog.conf>



# Log Files Management

- Log files can grow large quickly, especially with busy services, e.g., email, web, and DNS servers
- They may fill up the disk, degrading system performance
- Normally one uses a separate partition for busiest log files
  - On Linux, it is a good choice to have `/var` or `/var/log` occupy a separate partition on the disk

## Logs *not* to manage

Logs are text files to which lines are written as interesting events occur. But some logs are different

- **wtmp**: records of users' logins / logouts, system reboot and shutting down. Binary format. Use “**last**” command to decode
- **lastlog**: similar to above. Only records last login for each user.
- **utmp**: keeps a record of each user that is currently logged in. Maybe inaccurate if a shell is killed inappropriately
- You may read the man pages of each for more information

# Vendor specific log file locations

- Vendors may have their log files all over the disk. Check daemons' config files and syslog configuration files to find them

## Linux “logrotate” tool

- Linux logs are usually clearly named and consistently stored in `/var/log`
- Linux distributions also include a log management tool “logrotate”. It rotates, truncates, manages logs
- New software can add a config file to `/etc/logrotate.d` directory, to set up a management strategy for their logs, as part of their installation procedure.

# Syslog: the system event logger

- Liberate programmers from tedious mechanics of writing log files
- Put administrators in control of logging rather than letting every program make up its own logging policy, such as what information to keep and where it is stored
- Let you sort messages by importance and source, also route messages to a variety of destinations: **log files, users' terminals, other machines' syslogd**
  - The last one can centralize logging on a network

# Syslog Architecture

Three parts:

- **syslogd**: the logging daemon, its config file `/etc/syslog.conf`
  - **openlog** et al., library routines that submit msgs to syslogd
  - **logger**: a user command that submits log entries from the shell
- 
- Syslogd is started at boot time and runs continuously
  - Programs write log entries using the library calls
  - One can submit an entry using command “logger”
    - `logger -p local7.warning “a warning message”`

# Configuring Syslogd

- `/etc/syslog.conf` file, called `/etc/rsyslog.conf` in CentOS 6
  - It is a text file with simple format
  - ‘#’ starts comment lines, which are ignored
  - The basic format: `Selector<tab>action`
  - Can have one or more tabs
  - E.g., “`mail.info<tab>/var/log/maillog`”  
causes messages from the email system to be saved in `/var/log/maillog` file

# Syslog Selectors

- Selectors identify the program sending the log message , and the message's severity level,
- Selectors syntax *facility.level*
  - Both facility names and severity levels must be from a short list of defined values
  - Facilities are defined for the kernel, for common utilities, for locally written program, and for others named “user”
  - Also use special keywords: *\** means all, *none* means nothing, *comma* to separate multiple facilities, *;* to separate multiple selectors
  - *Facility* names: auth, cron, daemon, ftp, kern, local0-7, lpr, mail, news, ...  
*Severity levels* (descending severity): emerg, alert, crit, err, warning, notice, info, debug

# Syslog Actions

Syslog produces time stamp messages.

- *Filename*: appends the message to a file on the local machine
- *@hostname*: forwards the message to the syslogd on hostname
- */fifoname*: writes the message to the named pipe
- *User1, user2*: write the message to the screens of users if they are logged in
- *\**: write the message to all users currently logged in
- *-*: means no filesystem syncing after writing each log entry, this helps with performance, may miss some log upon crash



# Linux “logrotate” tool

- “logrotate” rotates, truncates, manages logs
- The logrotate config file is </etc/logrotate.conf>
- logrotate is normally run out of cron once a day
- Example logrotate options:
  - Compress all noncurrent versions of the log files
  - Rotate log files daily, weekly, or monthly
  - Emails error notification to a specified email address
  - Specify script to run after log is rotated
  - Include  $n$  versions of log

# Condensing Logs

- Syslog great for sorting and routing log messages, at the end a bunch of log files are created
- Tools can scan log entries, match against a database of patterns of log messages, and find the important messages
- Example **log postprocessor** tools: *swatch*, *logcheck*, *Splunk*, *SEC (Simple Event Correlator)* etc.
  - *swatch*: 'simple watchdog' to monitor log files from syslog and others

# Important Checking

Always check for important items, including:

- Most security-related messages need prompt review
  - Failed login, su, sudo attempts. Someone may forget passwords, but also want to prevent potential break-ins
- Messages about disks that have filled up
  - Full disks often bring useful work to a standstill
- Events that repeated many times

# Logging Policies

- Logs are critical to security incident handling
- Ask the following when designing logging policies
  - How many systems and apps will be included?
  - What type of storage infrastructure is available?
  - How long must logs be retained?
  - What types of events are important?
- Record the following:
  - user name or ID, event success or failure, source address, data and time, sensitive data changed, event details

# Log Centralization

- If site has >20 servers, consider centralized log collection and analysis. Reasons:
  - Simplified storage, automated analysis and alerting, improves attack visibility
- Storage strategy:
  - E.g., 30days on RAID array, 1 year on SAN, and 3 years on tape archives
- Access only to high-level sysadmins, access to central logs should be logged
- Small sites: rotate logs, regular archives