Remote Logging with Rsyslog

Or, How I Learned to Start Worrying and Love the Panopticon

Paul Nijjar

Kitchener-Waterloo Linux User Group

August 10, 2009

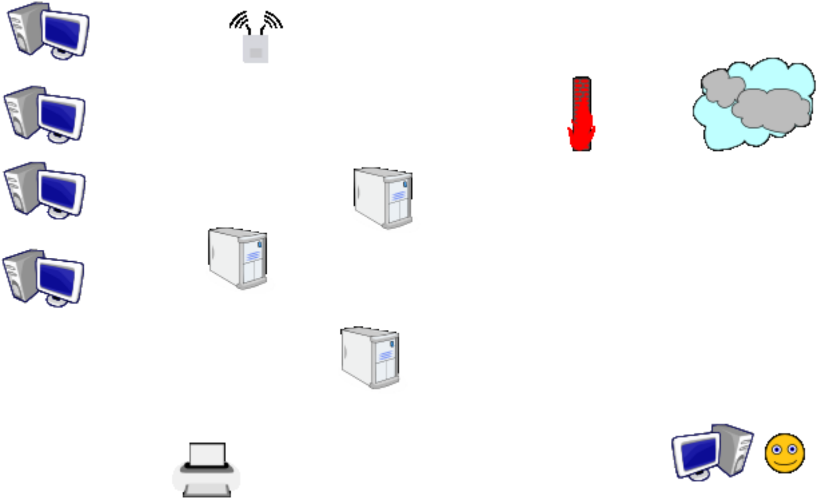

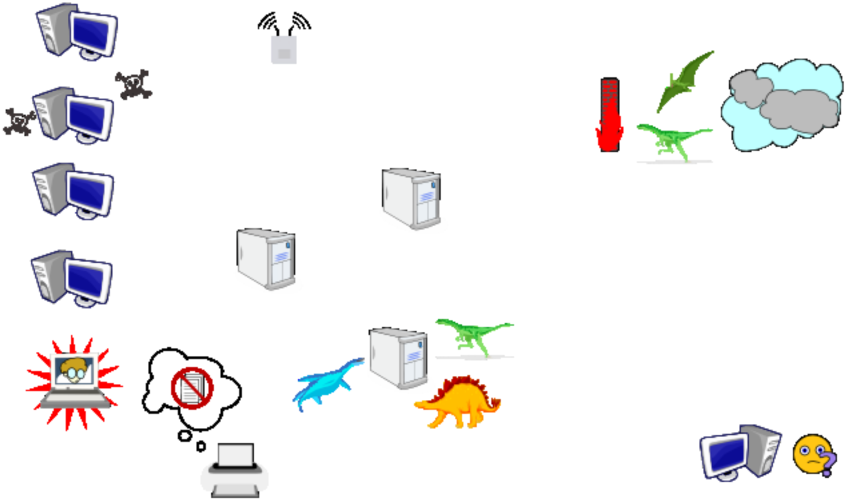

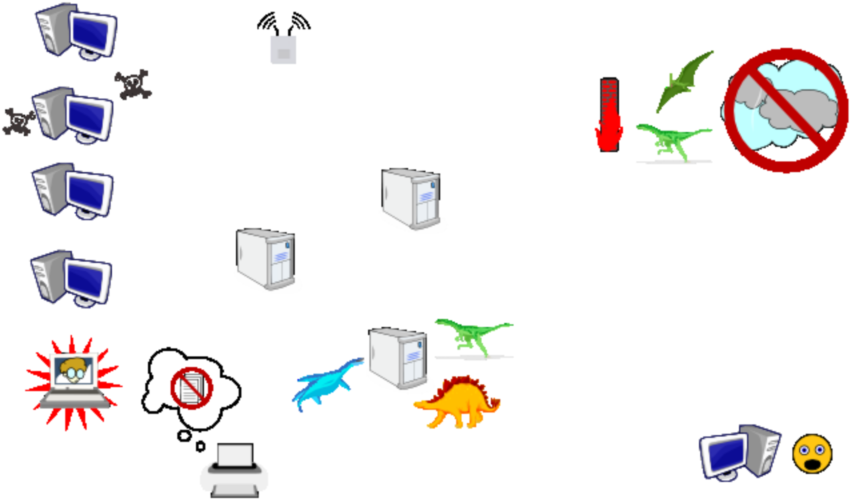

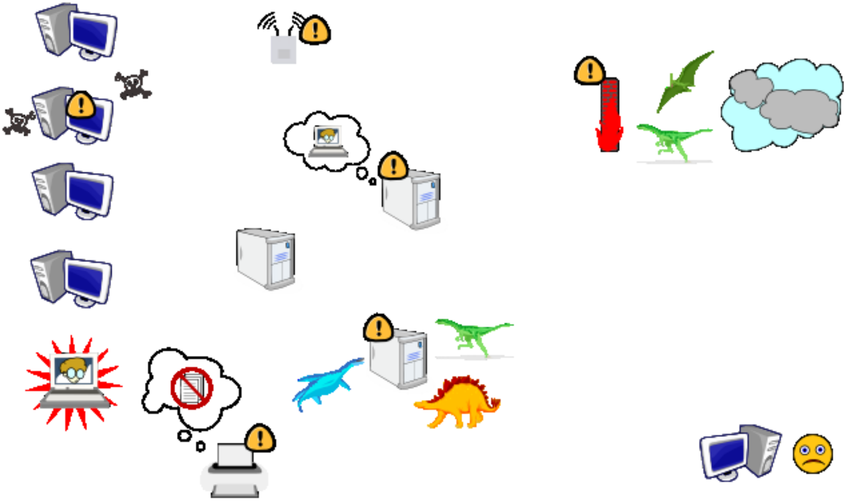

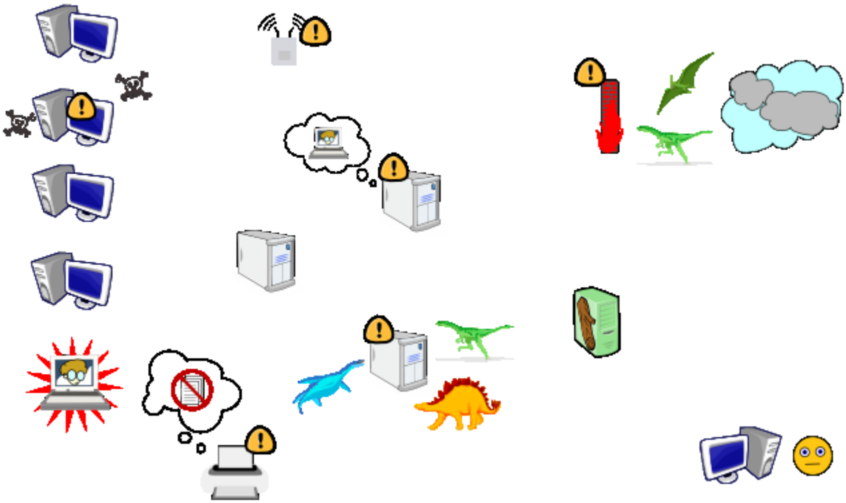

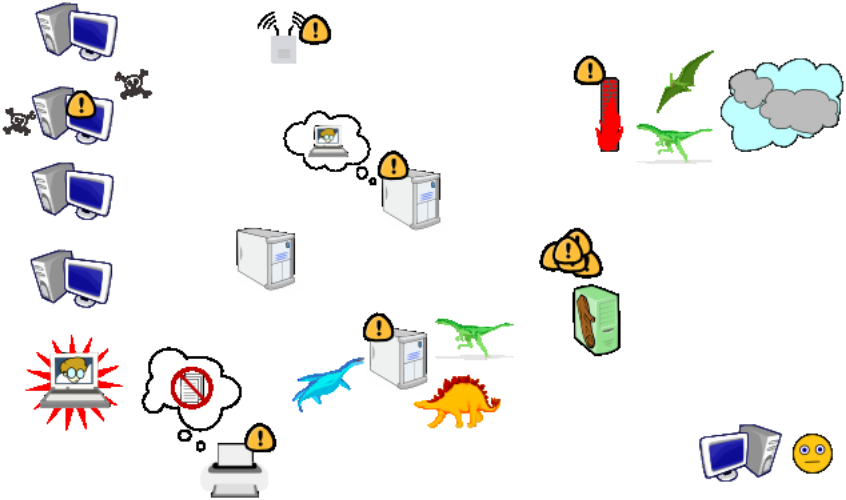

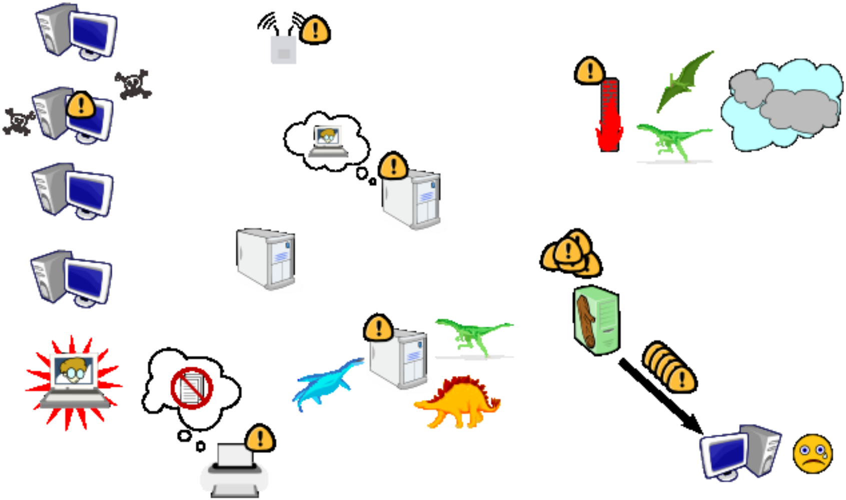

Goals

Centralize Logging: Look in one place, using one set of tools.

Archive Logs: Keep logs around for at least a year.

Generate Alerts: Tell me when something goes wrong.

Identify Trends: Tell me what “business as usual” looks like.

The last two of these goals are still works in progress.

Another goal: do this on the cheap, preferably with FLOSS.

Rsyslog

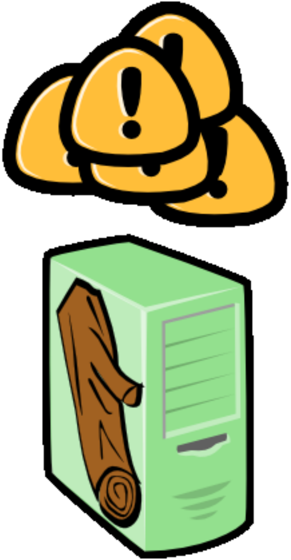

About Syslog

Syslogd is a logging interface used by many Linux programs to write log files. It is responsible for:

- Many of the files in `/var/log`: `messages`, `debug`, `syslog`, etc.
- Messages sent to the system console.
- Messages forwarded to other systems.
- Emergency log messages printed on everybody's screens

About Rsyslog

Rsyslog is a drop-in replacement for regular syslog. It adds a bunch of features:

- Better security controls
- More filtering options/syntax
- More reliable transport mechanisms
- Writing to databases

Rsyslog is now the default syslogging daemon for Fedora and Debian.

Configuring Rsyslog

1. Enable remote logging
2. Write templates for filenames and log formats
3. Filter messages from different hosts to different files
4. Rotate and archive files using `logrotate`
5. Debug the collection process

Config Files

In Debian, configuration is done in `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf`

Order matters, so I prepend configuration snippets with numbers:

- `/etc/rsyslog.d/00-AllowedHosts.conf`
- `/etc/rsyslog.d/40-Windows-Servers.conf`
- `/etc/rsyslog.d/99-EverythingElse.conf`

In general rules need to begin in the first column (no spaces) and they should be on one line.

Enabling Remote Logging

In `/etc/rsyslog.conf`, uncomment the following lines:

```
$ModLoad imudp
$UDPServerRun 514
```

UDP on port 514 is the standard syslog port.

You may need to open this port on your firewalls if you are logging from remote subnets/devices.

Allowing Remote Hosts to Syslog

In `/etc/rsyslog.d/00-AllowedHosts.conf`, allow some hosts. You can specify IP addresses, subnets, or hostnames:

```
# One server or router
$AllowedSender UDP, 192.168.1.4

# Everything in a subnet
$AllowedSender UDP, 192.168.2.0/24

# Everything (claiming to be) from Microsoft
$AllowedSender UDP, *.microsoft.com
```

Text Log Goals

My goal: put one or two logfiles per host in `/var/log/remote-logs/`

I don’t want to touch the local logging (e.g. `/var/log/messages`) at all.

I want to keep the logs for at least a year, and archive them in `/var/log/remote-logs/oldlogs/`

Log Message Properties

Every log message comes with some attributes called properties. Here are a few useful ones:

msg Message body.

rawmsg The message text as sent over the wire.

HOSTNAME The host that generated the message.

FROMHOST The host that last relayed this message.

syslogtag The service that reported the message. e.g. `kernel:`, `security[success]`

PRI The facility.priority of the message. e.g. `mail.debug`

There are others that can be useful for auditing, such as `timereported`, `syslogfacility-text`...

HOSTNAME va FROMHOST

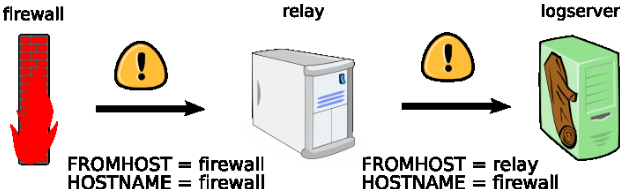

Templates

Templates are formatted strings. They can be used to name destination files and rewrite the format of messages that go to the syslog server.

```
$template BoringServerLog,
  "/var/log/remote-logs/%HOSTNAME%-boring.log"
```

e.g.

```
Aug 7 04:29:49 localhost su[2569]:
pam_authenticate:  Authentication failure
```

```
$template TraditionalFormat,
  "%timegenerated% %HOSTNAME% %syslogtag%%msg%\n"
```

Message Destinations

You can send messages to files (with an optional format):

```
*.*             /var/log/everything.log
*.debug         ?DebugLog
mail.*          ?MailLog;MailFormat
```

To stop processing messages send them to the ˜ destination:

```
*.*             ~
```

Filtering Messages With Selection Rules

Rsyslog provides four mechanisms for filtering messages into files:

BSD blocks Filter messages by hostname or program name

Traditional Filter by severity and facility

Property based Look at the message properties

Expression based If-then statements

I could get the first three of these to work.

Note that you cannot mix these methods on one line (but you can put other rules inside a BSD block)

BSD blocks

Specify a hostname for which all following rules will apply:

```
+mailserver

*.*     /var/log/remote-logs/mailserver.log
*.*     ~
```

You can make rules for all but a certain host

```
-mailserver

*.*     /var/log/remote-logs/allbutmail.log
```

BSD Blocks

You can unset the code block afterwards to allow all hosts.

```
+*

*.*     /var/log/everybody.log
```

There is also syntax that allows you to make blocks based on program name:

```
!sudo

*.*     /var/log/sudostuff.log
```

Traditional Selectors

This is the standard `facility.priority` filtering from regular syslog.

Some facilities: `auth, authpriv, cron, daemon, local0, local1, local7, user`

Some priorities: `debug, info, notice, warning, err, crit, alert, emerg` .

By default specifying a priority includes messages from higher priorities to the same file.

Selector Examples

```
*.*             /var/log/everything.log

# daemon messages of priority err to emerg
daemon.err      /var/log/daemon-warning.log

# Only messages of priority crit
*.=crit         /var/log/critical.log

# Emergencies get printed on everybody’s screen
*.emerg         *
```

Property Based Filters

These allow you to filter based on message properties. They begin with a colon.

```
:msg, contains, "RGFW-OUT: ACCEPT (ICMP type 8"
  /var/log/remote-logs/pubrouter-stupid.log
```

Property-based filters are slower than traditional ones, but I used them a lot.

Property Operators

The following operators are defined:

isequal Does the property match exactly?

contains Does the property contain a string?

startswith Does the property start with a certain string?

regex Does the property match a given regular expression?

Property Filter Examples

If `HOSTNAME` is not defined I often filtered like this:

```
:FROMHOST, isequal, "192.168.1.20"
   /var/log/remote-logs/pubrouter.log
```

Some messages in my router were of the form `192.168.1.42:28268 -> 192.161.1.3:443` for DNS lookups.

```
:msg, regex, ".*:443$" ?BoringDNSLog
```

If-Then Expression Filters

You are supposed to be able to use expressions filters like this:

```
if $FROMHOST isequal ’192.168.1.20’
  and $msg contains ’RGFW-OUT’
  /var/log/router-out.log
```

I could never get these to work, but maybe I am just dumb. As Rsyslog matures this is supposed to get more powerful.

Putting It Together

Some of `/etc/rsyslog.d/40-winservers.conf`

```
+dc1

:syslogtag, startswith, "DHCP"
  /var/log/remote-logs/dc1-dhcp.log
:syslogtag, startswith, "DHCP"  ~

*.*     /var/log/remote-logs/dc1.log
*.*     ~
```

Another Approach: MySQL

If you install the `rsyslog-mysql` package, you can write logs to a MySQL database.

Caution: On Debian, this package creates an `rsyslog` database user that is more powerful than it needs to be.

The package puts a file called `mysql.conf` in `/etc/rsyslog.d/`, which I copied to a file called `07-mysql.conf` .

07-mysql.conf

```
$ModLoad ommysql

# <dbserver>,<dbname>,<dbuser>,<dbpass>;<template>
*.* :ommysql:localhost,Syslog,rsyslog,dbpassword
```

The template is optional – there is a default schema and template used.

phpLogCon

You can download a PHP frontend to the Rsyslog MySQL database called from `http://www.phplogcon.org`

Installation is manual but pretty easy: untar scripts into `/var/www` and run a configuration script.

Dependencies: `rsyslog-mysql`, `php5-mysql`, `php5-gd`, `libapache2-mod-php5`

phpLogCon interface

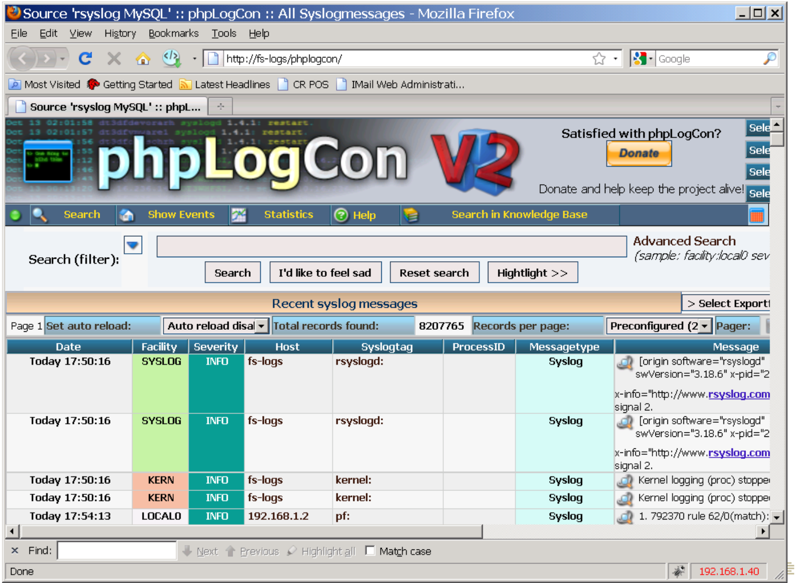

Page 1 | Set auto reload: | Auto reload disal | Total records found: | 8207765 | Records per page: | Preconfigured (2 | Pager:

Date	Facility	Severity	Host	Syslogtag	ProcessID	Messagetype	Message
Today 17:50:16	SYSLOG	INFO	fs-logs	rsyslogd:		Syslog	[origin software="rsyslogd" swVersion="3.18.6" x-pid="2 x-info="http://www.rsyslog.com signal 2.
Today 17:50:16	SYSLOG	INFO	fs-logs	rsyslogd:		Syslog	[origin software="rsyslogd" swVersion="3.18.6" x-pid="2 x-info="http://www.rsyslog.com signal 2.
Today 17:50:16	KERN	INFO	fs-logs	kernel:		Syslog	Kernel logging (proc) stoppe
Today 17:50:16	KERN	INFO	fs-logs	kernel:		Syslog	Kernel logging (proc) stoppe
Today 17:54:13	LOCAL0	INFO	192.168.1.2	pf:		Syslog	1. 792370 rule 62/0(match):

phpLogCon Messages

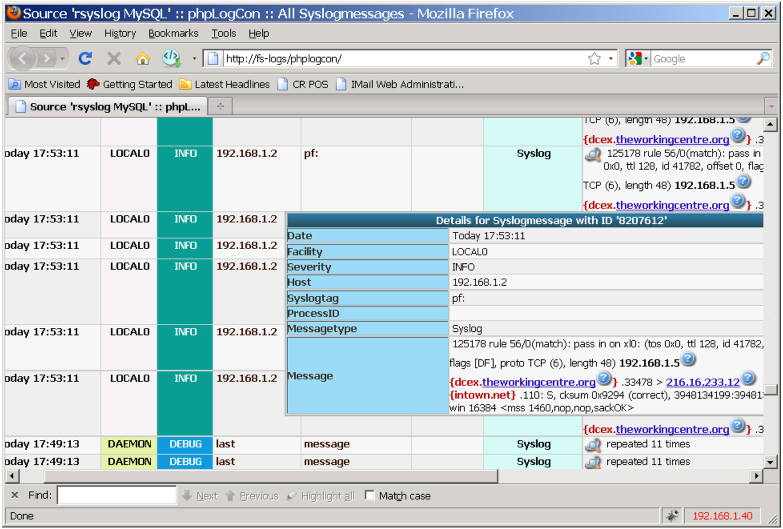
Source 'rsyslog MySQL' :: phpLogCon :: All Syslogmessages - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://fs-logs/phplogcon/
Google
Most Visited
Getting Started
Latest Headlines
CR POS
IMail Web Administrati...
Source 'rsyslog MySQL' :: phpL...
Find:
Next
Previous
Highlight all
□ Match case
Done
192.168.1.40

							TCP (6), length 48) **192.168.1.5** {dcex.theworkingcentre.org} .3
oday 17:53:11	LOCAL0	INFO	192.168.1.2	pf:		Syslog	125178 rule 56/0(match): pass in 0x0, ttl 128, id 41782, offset 0, flag TCP (6), length 48) **192.168.1.5** {dcex.theworkingcentre.org} .3
oday 17:53:11	LOCAL0	INFO	192.168.1.2				
oday 17:53:11	LOCAL0	INFO	192.168.1.2				
oday 17:53:11	LOCAL0	INFO	192.168.1.2				
oday 17:53:11	LOCAL0	INFO	192.168.1.2				
oday 17:53:11	LOCAL0	INFO	192.168.1.2				{dcex.theworkingcentre.org} .3
oday 17:49:13	DAEMON	DEBUG	last	message		Syslog	repeated 11 times
oday 17:49:13	DAEMON	DEBUG	last	message		Syslog	repeated 11 times

Details for Syslogmessage with ID '8207612'

Date	Today 17:53:11
Facility	LOCAL0
Severity	INFO
Host	192.168.1.2
Syslogtag	pf:
ProcessID	
Messagetype	Syslog
Message	125178 rule 56/0(match): pass in on xl0: (tos 0x0, ttl 128, id 41782, flags [DF], proto TCP (6), length 48) **192.168.1.5** {dcex.theworkingcentre.org} .33478 > 216.16.233.12 {intown.net} .110: S, cksum 0x9294 (correct), 3948134199:39481 win 16384 <mss 1460,nop,nop,sackOK>

phpLogCon Host Graph

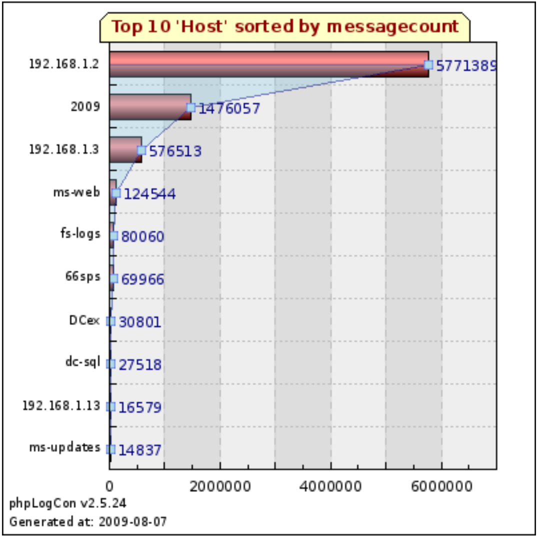

phpLogCon Severity Graph

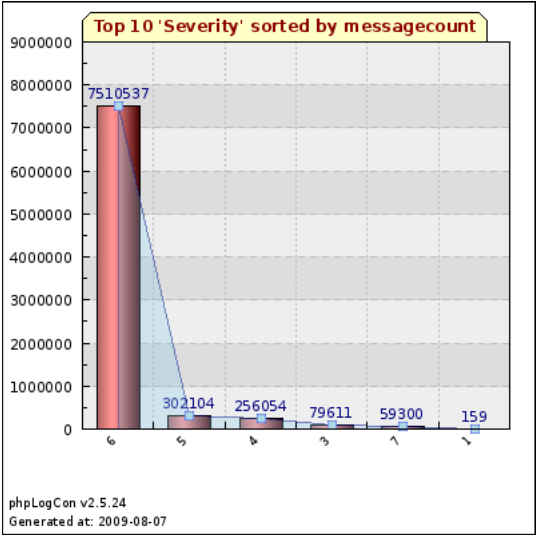

phpLogCon SyslogTag Graph

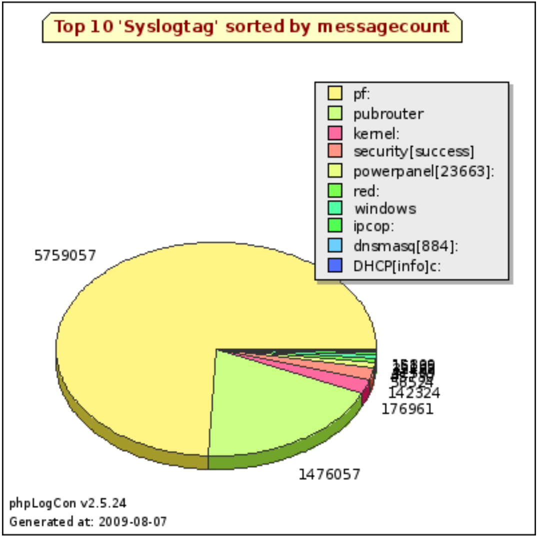

Archiving Logfiles

To archive logfiles I had to manually edit `/etc/logrotate.conf`. Most of it is pretty standard.

```
/var/log/remote-logs/*.log
{
  rotate 60
  weekly
  missingok
  notifempty
  compress
  delaycompress
```

This says: keep 60 weeks of logs. Compress old files, but wait a week before doing so. Don't archive empty files and don't complain about them.

```
  sharedscripts
  postrotate
    invoke-rc.d rsyslog reload > /dev/null
  endscript
  olddir /var/log/remote-logs/oldlogs
}
```

This says: restart `rsyslog` once after moving all files. Put the files in the `oldlogs` directory.

Debugging

Debugging can be hideous. Here are some tools to make it easier.

- Listing logs by update time
- DEBUG template
- Using logger to send messages locally
- Rsyslog in verbose mode
- Wireshark/TCPDump

Listing logs by update time

This is suprisingly handy to see if a particular host has been writing files recently. It sorts files by modification time.

```
ls -ltc
```

DEBUG template

In `rsyslog.d/05-DebugTemplate.conf` add the following template (given in the documentation):

```
$template DEBUG,"Debug line with all properties:
\nFROMHOST: ’%FROMHOST%’, HOSTNAME: ’%HOSTNAME%’,
PRI: %PRI%,\nsyslogtag ’%syslogtag%’,
programname: ’%programname%’,
APP-NAME: ’%APP-NAME%’, PROCID: ’%PROCID%’,
MSGID: ’%MSGID%’,\nTIMESTAMP: ’%TIMESTAMP%’,
STRUCTURED-DATA: ’%STRUCTURED-DATA%’,
\nmsg: ’%msg%’
\nescaped msg: ’%msg:::drop-cc%’
\nrawmsg: ’%rawmsg%’\n\n"
```

Use DEBUG template

Now in `/rsyslog.d/70-EverythingElse.conf` log every remote message that has not been logged already:

```
-logserver

$template RemoteHostLog,
  "/var/log/remote-logs/uncaught.log"

*.*     ?RemoteHostLog;DEBUG
*.*     ~
```

You can also activate this for particular hosts, or for hosts that do not have a `HOSTNAME` defined.

Using logger to send messages locally

You can use the `logger` command to write syslog messages manually:

```
# Send with priority user.info
logger ’I hate test messages!’

logger -p kern.emerg ’Everything is broken!’
```

Rsyslog in debug mode

This will produce a HUGE amount of information. It can be useful in checking whether your messages are getting to the daemon.

```
/etc/init.d/rsyslog stop
script /tmp/output.txt

rsyslogd -c3 -d
<ctrl>+C

exit
/etc/init.d/rsyslog start
```

Wireshark and tcpdump

This is useful to see whether messages are getting to the syslog server. Use the following filter to see what is coming in on UDP port 514:

```
udp.port == 514
```

The equivalent filter for `tcpdump` is:

```
tcpdump udp port 514
```

Sending Logs from Computers and Devices

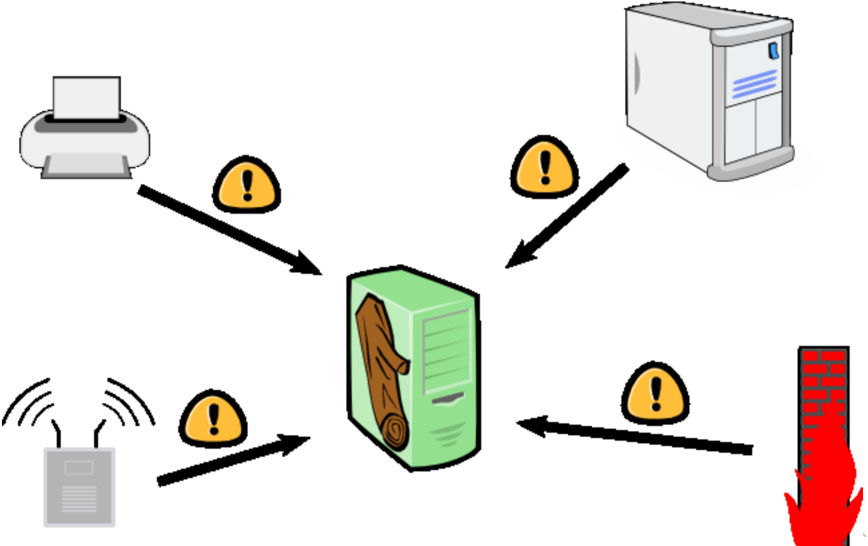

Sending logs from UNIX/Linux

In the syslogd.conf of the client, add the following line before any log messages are thrown away:

```
*.*     @192.168.1.40
```

This forwards messages using UDP over the default port. Many sysloggers support TCP as well (with @@).

Your client does not need to run rsyslog for this to work. Most sysloggers will work.

Sending logs from devices

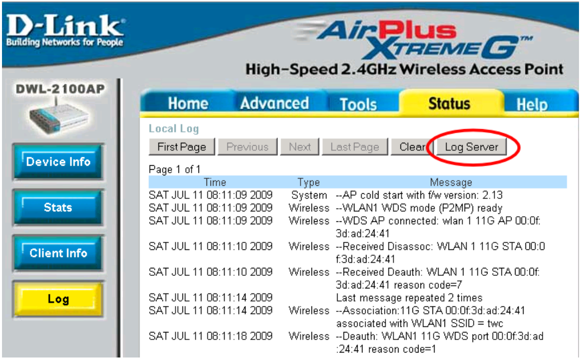

```
Page 1 of 1
          Time               Type                  Message
SAT JUL 11 08:11:09 2009    System   --AP cold start with f/w version: 2.13
SAT JUL 11 08:11:09 2009   Wireless  --WLAN1 WDS mode (P2MP) ready
SAT JUL 11 08:11:09 2009   Wireless  --WDS AP connected: wlan 1 11G AP 00:0f:
                                     3d:ad:24:41
SAT JUL 11 08:11:10 2009   Wireless  --Received Disassoc: WLAN 1 11G STA 00:0
                                     f:3d:ad:24:41
SAT JUL 11 08:11:10 2009   Wireless  --Received Deauth: WLAN 1 11G STA 00:0f:
                                     3d:ad:24:41 reason code=7
SAT JUL 11 08:11:14 2009             Last message repeated 2 times
SAT JUL 11 08:11:14 2009   Wireless  --Association:11G STA 00:0f:3d:ad:24:41
                                     associated with WLAN1 SSID = twc
SAT JUL 11 08:11:18 2009   Wireless  --Deauth: WLAN1 11G WDS port 00:0f:3d:ad
                                     :24:41 reason code=1
```

Enabling Devices

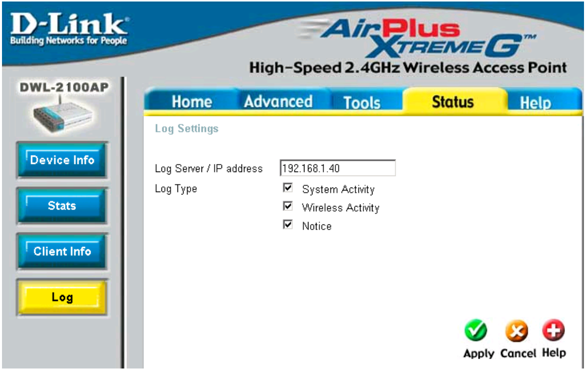

Log Settings

Log Server / IP address 192.168.1.40

Log Type ☑ System Activity
☑ Wireless Activity
☑ Notice

WARNING

The following slides contain depictions of proprietary software use and may not be suitable for all viewers. Viewer discretion is advised.

Windows AUGH

Naturally, Windows does not speak syslog format natively. However, there are tools to convert Windows event logs to syslog format.

Windows Vista/2008 introduced an XML format `.evtx` which I don't care about (yet).

SyslogAgent

This is commercial software released under the GPL. Get it from `http://syslogserver.com/syslogagent.html`

This runs as a system service.

There are a few other syslog agents available. (The Rsyslog guy makes a proprietary one.)

I found SysLogAgent lightweight, easy to install, and good enough for my purposes.

SysLogAgent main screen

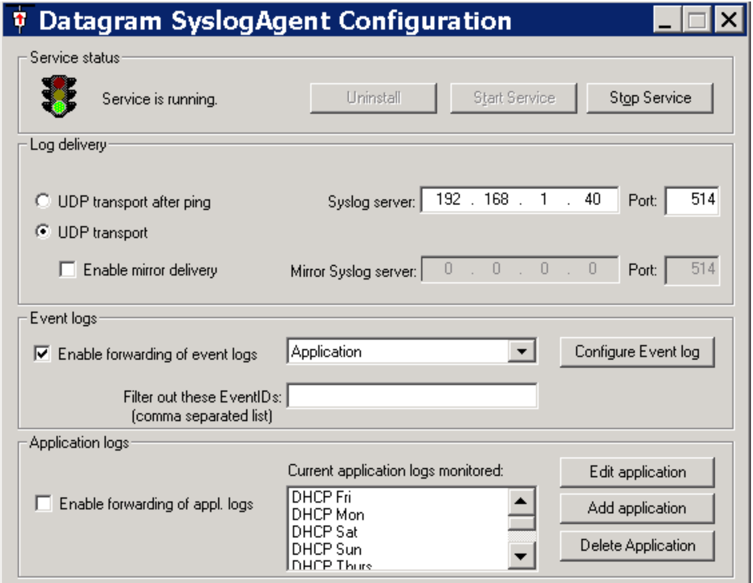

Datagram SyslogAgent Configuration

Service status

Service is running.

Log delivery

○ UDP transport after ping

● UDP transport

□ Enable mirror delivery

Syslog server: 192 . 168 . 1 . 40 Port: 514

Mirror Syslog server: 0 . 0 . 0 . 0 Port: 514

Event logs

☑ Enable forwarding of event logs

Application

Filter out these EventIDs:
(comma separated list)

Application logs

□ Enable forwarding of appl. logs

Current application logs monitored:

DHCP Fri
DHCP Mon
DHCP Sat
DHCP Sun
DHCP Thurs

SysLogAgent: Specifying Messages to Send

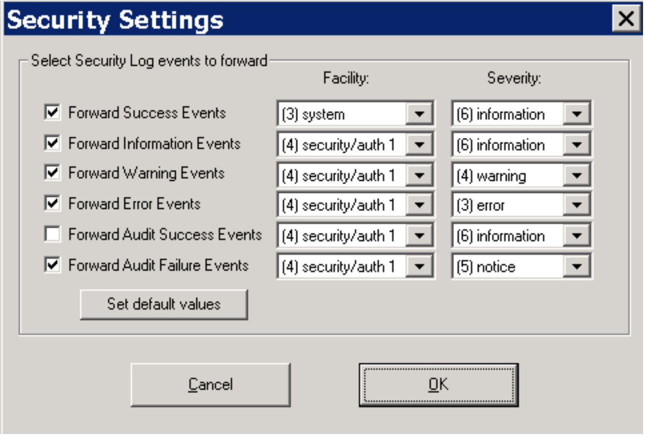

Security Settings

Select Security Log events to forward

	Facility:	Severity:
☑ Forward Success Events	(3) system	(6) information
☑ Forward Information Events	(4) security/auth 1	(6) information
☑ Forward Warning Events	(4) security/auth 1	(4) warning
☑ Forward Error Events	(4) security/auth 1	(3) error
☐ Forward Audit Success Events	(4) security/auth 1	(6) information
☑ Forward Audit Failure Events	(4) security/auth 1	(5) notice

Generating Windows Events

There is a commandline interface to generate Windows System Log events called `eventcreate.exe`

```
eventcreate /t ERROR /id 666
  /d "Our stock price is falling!"
```

A Story: MS DHCP Logs

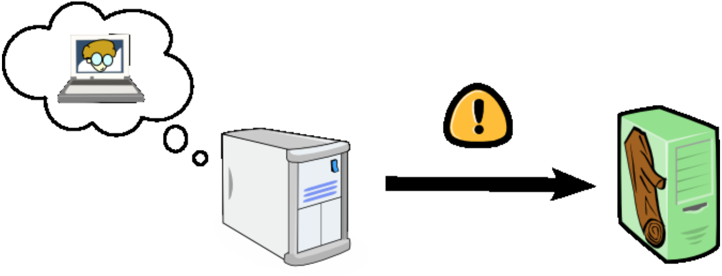

A Sad Story

Microsoft’s DHCP server can write out pretty good logs.

Naturally, they don’t show up as events in Event Viewer.

Instead, they are textfiles in `c:\windows\system32\dhcp\`

How can we get them into the syslog server?

More Sadness

The first 30 lines in every logfile are purely informational:

```
                     Microsoft DHCP Service Activity Log


Event ID  Meaning
00        The log was started.
01        The log was stopped.
02        The log was temporarily paused due to low disk space
...
32        DNS update successful
50+       Codes above 50 are used for Rogue Server Detection i

ID,Date,Time,Description,IP Address,Host Name,MAC Address
24,08/06/09,00:00:57,Database Cleanup Begin,,,,
30,08/06/09,00:00:57,DNS Update Request,85.1.168.192,EM14temp.
25,08/06/09,00:00:57,0 leases expired and 0 leases deleted,,,,
```

Even Sadder

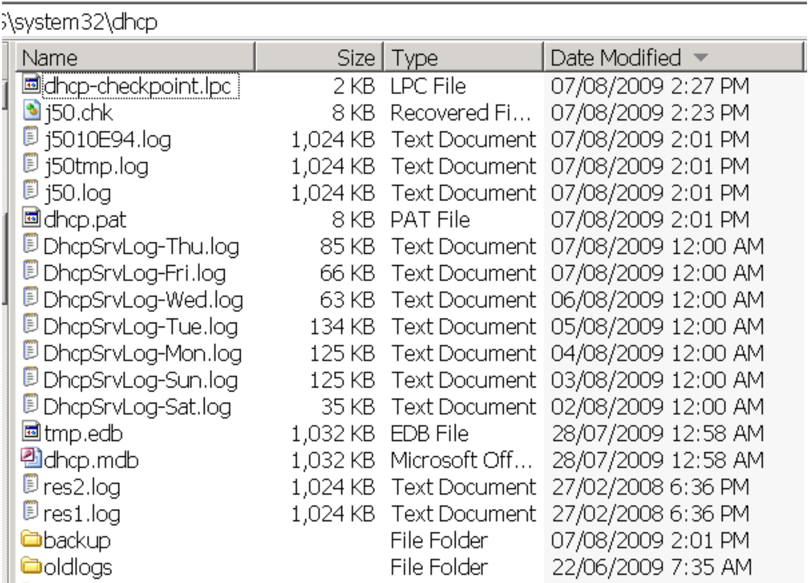

\system32\dhcp

Name	Size	Type	Date Modified
dhcp-checkpoint.lpc	2 KB	LPC File	07/08/2009 2:27 PM
j50.chk	8 KB	Recovered Fi...	07/08/2009 2:23 PM
j5010E94.log	1,024 KB	Text Document	07/08/2009 2:01 PM
j50tmp.log	1,024 KB	Text Document	07/08/2009 2:01 PM
j50.log	1,024 KB	Text Document	07/08/2009 2:01 PM
dhcp.pat	8 KB	PAT File	07/08/2009 2:01 PM
DhcpSrvLog-Thu.log	85 KB	Text Document	07/08/2009 12:00 AM
DhcpSrvLog-Fri.log	66 KB	Text Document	07/08/2009 12:00 AM
DhcpSrvLog-Wed.log	63 KB	Text Document	06/08/2009 12:00 AM
DhcpSrvLog-Tue.log	134 KB	Text Document	05/08/2009 12:00 AM
DhcpSrvLog-Mon.log	125 KB	Text Document	04/08/2009 12:00 AM
DhcpSrvLog-Sun.log	125 KB	Text Document	03/08/2009 12:00 AM
DhcpSrvLog-Sat.log	35 KB	Text Document	02/08/2009 12:00 AM
tmp.edb	1,032 KB	EDB File	28/07/2009 12:58 AM
dhcp.mdb	1,032 KB	Microsoft Off...	28/07/2009 12:58 AM
res2.log	1,024 KB	Text Document	27/02/2008 6:36 PM
res1.log	1,024 KB	Text Document	27/02/2008 6:36 PM
backup		File Folder	07/08/2009 2:01 PM
oldlogs		File Folder	22/06/2009 7:35 AM

Logs everywhere!

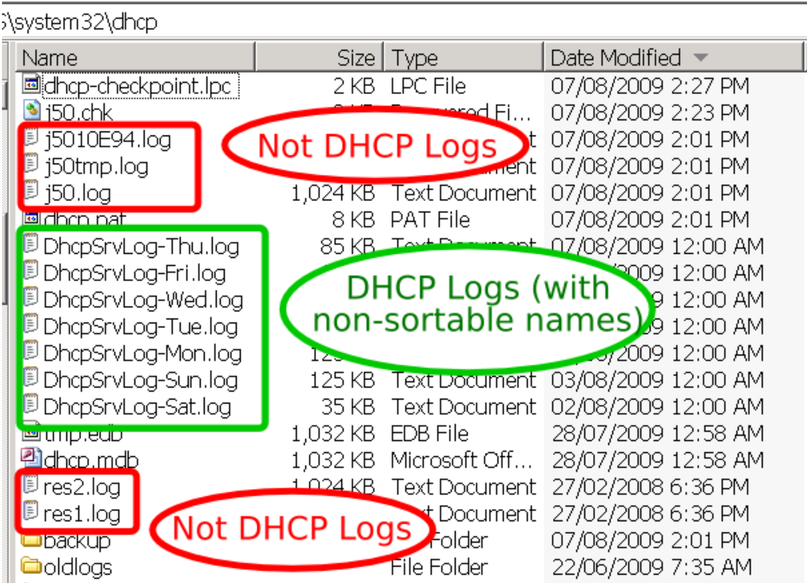

\system32\dhcp

Name	Size	Type	Date Modified
dhcp-checkpoint.lpc	2 KB	LPC File	07/08/2009 2:27 PM
j50.chk	[illegible]	ered Fi…	07/08/2009 2:23 PM
j5010E94.log	[illegible]	t	07/08/2009 2:01 PM
j50tmp.log	[illegible]	ment	07/08/2009 2:01 PM
j50.log	1,024 KB	Text Document	07/08/2009 2:01 PM
dhcp.pat	8 KB	PAT File	07/08/2009 2:01 PM
DhcpSrvLog-Thu.log	85 KB	Text Document	07/08/2009 12:00 AM
DhcpSrvLog-Fri.log	[illegible]	[illegible]	009 12:00 AM
DhcpSrvLog-Wed.log	[illegible]	[illegible]	9 12:00 AM
DhcpSrvLog-Tue.log	[illegible]	[illegible]	9 12:00 AM
DhcpSrvLog-Mon.log	12	[illegible]	/2009 12:00 AM
DhcpSrvLog-Sun.log	125 KB	Text Document	03/08/2009 12:00 AM
DhcpSrvLog-Sat.log	35 KB	Text Document	02/08/2009 12:00 AM
tmp.edb	1,032 KB	EDB File	28/07/2009 12:58 AM
dhcp.mdb	1,032 KB	Microsoft Off…	28/07/2009 12:58 AM
res2.log	1,024 KB	Text Document	27/02/2008 6:36 PM
res1.log	[illegible]	t Document	27/02/2008 6:36 PM
backup		Folder	07/08/2009 2:01 PM
oldlogs		File Folder	22/06/2009 7:35 AM

Not DHCP Logs

DHCP Logs (with non-sortable names)

Not DHCP Logs

SysLogAgent!

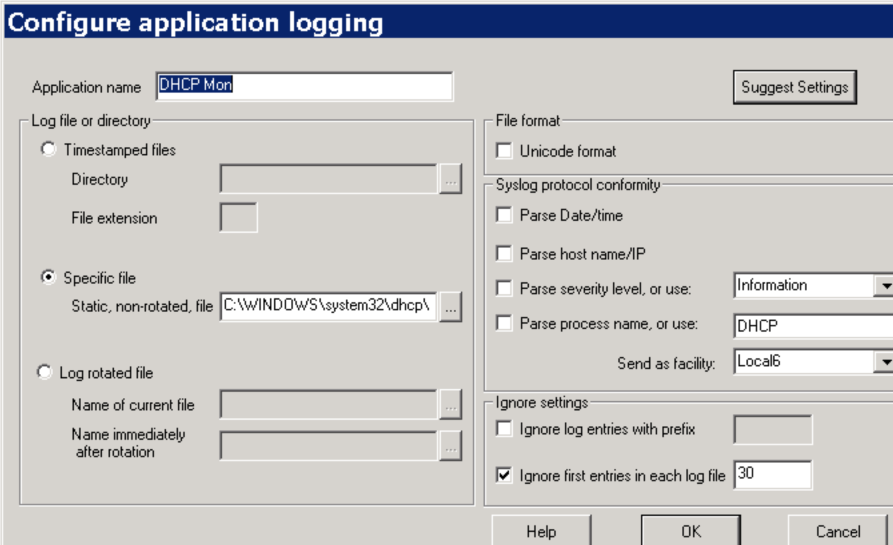

Configure application logging

Application name: DHCP Mon

Suggest Settings

Log file or directory

○ Timestamped files

Directory:

File extension:

● Specific file

Static, non-rotated, file: C:\WINDOWS\system32\dhcp\

○ Log rotated file

Name of current file:

Name immediately after rotation:

File format

□ Unicode format

Syslog protocol conformity

□ Parse Date/time

□ Parse host name/IP

□ Parse severity level, or use: Information

□ Parse process name, or use: DHCP

Send as facility: Local6

Ignore settings

□ Ignore log entries with prefix

☑ Ignore first entries in each log file: 30

Help | OK | Cancel

Foiled!

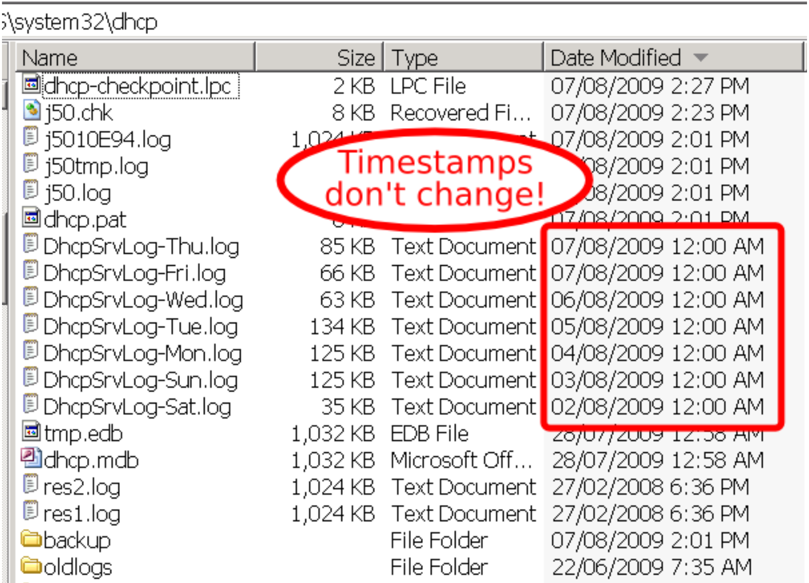

\system32\dhcp

Name	Size	Type	Date Modified
dhcp-checkpoint.lpc	2 KB	LPC File	07/08/2009 2:27 PM
j50.chk	8 KB	Recovered Fi...	07/08/2009 2:23 PM
j5010E94.log	1,024 KB	[illegible]	07/08/2009 2:01 PM
j50tmp.log	[illegible]	[illegible]	07/08/2009 2:01 PM
j50.log	[illegible]	[illegible]	07/08/2009 2:01 PM
dhcp.pat	[illegible]	[illegible]	07/08/2009 2:01 PM
DhcpSrvLog-Thu.log	85 KB	Text Document	07/08/2009 12:00 AM
DhcpSrvLog-Fri.log	66 KB	Text Document	07/08/2009 12:00 AM
DhcpSrvLog-Wed.log	63 KB	Text Document	06/08/2009 12:00 AM
DhcpSrvLog-Tue.log	134 KB	Text Document	05/08/2009 12:00 AM
DhcpSrvLog-Mon.log	125 KB	Text Document	04/08/2009 12:00 AM
DhcpSrvLog-Sun.log	125 KB	Text Document	03/08/2009 12:00 AM
DhcpSrvLog-Sat.log	35 KB	Text Document	02/08/2009 12:00 AM
tmp.edb	1,032 KB	EDB File	28/07/2009 12:58 AM
dhcp.mdb	1,032 KB	Microsoft Off...	28/07/2009 12:58 AM
res2.log	1,024 KB	Text Document	27/02/2008 6:36 PM
res1.log	1,024 KB	Text Document	27/02/2008 6:36 PM
backup		File Folder	07/08/2009 2:01 PM
oldlogs		File Folder	22/06/2009 7:35 AM

Timestamps don't change!

Log Parser 2.2

Microsoft has a freeware utility called Log Parser which can help. (Microsoft employees get frustrated by Windows too.)

It is a commandline “any-to-any” log converter with SQLesque syntax.

You can run it every minute with Task Scheduler

Log Parser Magic Syntax

```
LogParser.exe -i:TEXTLINE -iCheckPoint:check.lpc
  -o:SYSLOG -hostname:dc1 -processName:DHCP[info]
  "SELECT * INTO @192.168.1.40
  FROM DhcpSrvLog-*.log WHERE Index > 30"
```

-i:TEXTLINE Text file input

-iCheckPoint Remember the last location

-o:SYSLOG Syslog format output

INTO @192.168.1.40 Send to logserver

WHERE Index > 30 Skip first 30 lines

Tenshi

Who Watches the Logs?

My goals: be lazy but informed

- Get alerted when important things happen
- Get summaries of interesting log events
- Format the stuff so I will actually read it without feeling swamped.

Log Watchers

There are lots of them: `swatch`, `logwatch`, `sec`, `log2mail`, `logsentry`...

My arbitrary choice: `tenshi`

Tenshi concepts

Tenshi collects log messages into queues .

Identical messages are tallied in reports.

You can use masks to filter irrelevant information and make different messages appear identical to Tenshi.

Basic configuration

In `/etc/tenshi/tenshi.conf`

Specify logfiles to watch:

```
set logfile /var/log/auth.log
set logfile /var/log/remote-logs/dc1-dhcp.log
```

Limit report size from host (default is 800)

```
set limit 80
```

This says that a host may produce 80 lines of information per report.

Queues

Queues are used to sort messages and send them at different frequencies and in different ways. Syntax:

```
set queue <queue_name> <mail_from> <mail_to>
  <interval>
```

This queue will be flushed at most every two minutes. If there are no alerts it will do nothing. The subject will be “Log Alert!”

```
set queue important
  tenshi@localhost alerts@contoso.com
  [*/2 * * * *] Log alert!
```

More Queues

This queue goes out every Wednesday at 4:20pm with the default subject (which can be set in `tenshi.conf`)

```
set queue report tenshi@localhost
  list@contoso.com  [20 16 * * Wed]
```

This queue goes out immediately and is sent to a pager and a mailing list with the subject “Emergency!”

```
set queue emergency tenshi@localhost
  pager:37337@pager.com,alert@contoso.com
  [now] Emergency!
```

A builtin queue called `trash` is used to ignore messages entirely.

Tenshi rules

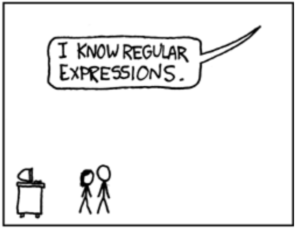

Rules for filtering messages go in `/etc/tenshi/includes-active/`

They are specified using regular expressions.

Like `rsyslog`, order matters.

Unlike `rsyslog` the first rule that applies “eats” the message.

Sample Rules

Context: Firewall messages look like this:

```
pf: 138214 rule 60/0(match): block in on
em0: (tos 0x0, ttl 118, id 49601, offset 0,
flags [DF], proto TCP (6), length 58)
209.73.191.147.1935 > 174.113.185.28.59609:
P, cksum 0x8198 (correct), 0:18(18) ack
1 win 65535
```

They all begin with `pf:`

I am largely interested in the IP addresses and ports:

```
209.73.191.147.1935 > 174.113.185.28.59609:
```

Code blocks

Apply the following rules onto to messages beginning with `pf:`

```
group ^pf:
```

which you end with

```
group_end
```

Report any message that comes from an address and port 6669

```
important \d+\.\d+\.\d+\.\d+\.6669[ :]
```

Mask out almost everything about firewall traffic that passes using the parentheses.

```
report (\d+) .+? pass in on .+?: (.+)
```

Sample output:

```
___ rule 54/0(match): pass in on xl0: ___
```

Down the slippery slope

Actual tenshi rules. They match things like:

```
DHCP[info]c:\WINDOWS\system32
\dhcp\DhcpSrvLog-Tue.log
78 8 10,08/04/09,09:29:04,Assign,
192.168.1.66,58tf-loftX.,00065BCAF8BD,
```

```
report ^DHCP\[.+?](.+?)
  \d\d,(\d\d\/\d\d\/\d\d,\d\d:\d\d:\d\d,)
  Renew

dhcp ^DHCP\[.+?](c:.+?\.log) .*?,Assign

important ^DHCP\[.+?](c:.+?\.log) .*?,Conflict
```

“I know regular expressions!”

> Some people, when confronted with a problem, think “I know, I’ll use regular expressions.” Now they have two problems.

–Jamie Zawinski, August 1997

If you disagree, you might check out the `logwatch-database` package.

Lessons Learned

- Collecting and archiving logs can be worthwhile.
- `rsyslog` offers lots of new features and flexibility over standard syslog.
- You can get syslog files from a lot of places (but the formatting is often wretched).
- Alerts for expected events work okay.
- I'm still unhappy with reporting. Regular expressions are not the right tool.

Thoughts and Future Work

Thought: Log reporting is like spam filtering.

Idea: Use database reporting to flag messages that I want reported always, and to report any brand new messages I have never seen.

Thank You!

- OpenClipArt and its many contributors for releasing beautiful images I can use for free
- NetDirect for the projector
- The Working Centre for not firing me even though I embezzled hardware and company time for this presentation
- Randall Munroe at `xkcd.com` and Jamie Zawinski for quotations
- The authors of `rsyslog`, `SysLogAgent`, `tenshi` and many other tools for giving me software to present about
- The LaTeX, `latex-beamer`, GIMP, and Inkscape people for giving me tools to make this presentation.

OpenClipArt Credits

In no particular order. I used real names when I could conveniently find them, handles on `openclipart.org` or the filenames otherwise.

- Nicubunu (smileys)
- Gerald G (log/campfire)
- nicolas (Wireless box)
- Luiz Araujo (alert)
- Linda Kim (crossbones)
- Francesco Rollandin (dinosaurs)
- Andrew Fitzsimon (laptop, printer)
- denco (firewall)
- Chris Goerner (fly)

More OpenClipArt Credits

Still in no particular order:

- glenn rolla (workstation)
- Nicolas cl (Internet cloud)
- Nicu Buculei (no symbol)
- teudimundo (server)
- Jarna Vasmaa (paper)
- mimooh (server)
- Feth Arezki (xbill)

The End