What is a Bro log?

Justin Azoff

Aug 26, 2014

What is a Bro log?

A Bro log is a stream of high level entries that correspond to network events.

- A file downloaded via HTTP
- An email sent using SMTP
- A login over SSH

Not log, but logs.

Bro does not have a single "alert" type log. Instead each kind of event stream has a dedicated file with it's own set of fields.

Why more than one file?

- The SMTP log has 'from' and 'subject' fields
- The HTTP log has 'method' and 'uri' fields
- The 'from' field would not make sense for HTTP, and 'uri' does not make sense for SMTP

How many log files are there?

By default, bro will output about two dozen log files, depending on what types of traffic it can see:

conn.log dhcp.log dns.log dpd.log files.log http.log intel.log known_certs.log known_hosts.log known_services.log modbus.log notice.log radius.log smtp.log snmp.log socks.log software.log ssh.log ssl.log syslog.log traceroute.log weird.log x509.log

Signal to noise ratio

The main way that log files can be categorized is by their size and signal to noise ratio. Some logs files are large and will contain entries that can be either benign or malicious. Other files are smaller and contain more actionable information.

- 24K known_services.log
- 28K software.log
- 68K notice.log
- 311M dns.log
- 856M conn.log

High signal log files

Inventory related log files

These log files are updated once per day and inventory your network

- known_hosts.log
- known_services.log
- known_certs.log
- software.log

Other high signal files

- notice.log - When bro detects something it thinks is exceptional it raises a notice.
- intel.log - Traffic that matches lists of known bad indicators is logged here.

Aside - Customizing log file contents.

Bro makes it easy to take a large log file and filter a subset of the entries to a smaller file with a higher signal to noise ratio.

Examples

- Filtering the http.log to http_exe.log
- Filtering the http.log to http_wget.log
- Filtering the http.log to http_java.log
- Filtering the conn.log to conn_cn.log
- Filtering the ssh.log to ssh_non_us.log

What exactly does a stream of events look like?

The short answer: A CSV file.

We can create some log files by starting Bro and running the unix command:

```
curl www.google.com
```

This will request the google home page, but not any of the associated javascript or image files.

Bro will write an entry in the http.log describing this event. The http.log contains 27 columns which can be a bit daunting. We can transpose the columns into rows to make this single line from http.log easier to understand

http.log transposed

Field	Type	Value
ts	time	1408828734.304076
uid	string	CZceY8wvnES5foJp4
id.orig_h	addr	192.168.43.222
id.orig_p	port	65032
id.resp_h	addr	74.125.226.50
id.resp_p	port	80
trans_depth	count	1
method	string	GET
host	string	www.google.com
uri	string	/
referrer	string	-

http.log transposed

Field	Type	Value
user_agent	string	curl/7.30.0
request_body_len	count	0
response_body_len	count	21232
status_code	count	200
status_msg	string	OK
info_code	count	-
info_msg	string	-
filename	string	-
tags	set[enum]	(empty)

http.log transposed

Field	Type	Value
username	string	-
password	string	-
proxied	set[string]	-
orig_fuids	vector[string]	-
orig_mime_types	vector[string]	-
resp_fuids	vector[string]	FvwPGj436gbcfXpCGf
resp_mime_types	vector[string]	text/html

Not just http.

This one HTTP download caused Bro to write entries to 6 log files:

- http.log has the above entry
- dns.log has an entry from the dns query for www.google.com
- files.log has an entry from the html file that was downloaded
- conn.log has an entry for both the dns an http connections
- known_hosts.log has an entry for 192.168.43.222
- software.log has an entry for an HTTP::BROWSER of curl/7.30.0 seen on 192.168.43.222

known_hosts.log transposed

Field	Type	Value
ts	time	1408828734.303825
host	addr	192.168.43.222

192.168.43.222 was seen for the first time at 1408828734.303825

software.log transposed (slightly edited)

Field	Type	Value
ts	time	1408828734.304076
host	addr	192.168.43.222
software_type	enum	HTTP::BROWSER
name	string	curl
version.major	count	7
version.minor	count	30
version.minor2	count	0
unparsed_version	string	curl/7.30.0

curl/7.30.0 was seen for the first time on 192.168.43.222 at 1408828734.304076