pfSense - Bug #7020

<Hostname> is omitted when sending logs on syslog

12/19/2016 09:44 AM - Idar Lund

Status:	Feedback	**Start date:**	12/19/2016
Priority:	Normal	**Due date:**	
Assignee:		**% Done:**	0%
Category:	Logging	**Estimated time:**	0.00 hour
Target version:			
Affected Version:		**Affected Architecture:**	

Description

When sending "filterlog" over syslog the standard defined in https://doc.pfsense.org/index.php/Filter_Log_Format_for_pfSense_2.2 (<Timestamp> <Hostname> filterlog: <CSV data>) is not followed.

According to https://www.ietf.org/rfc/rfc3164.txt, a message that are sent to a remote host should include a header. In 4.1.2, the HEADER is explained; The HEADER contains two fields called the TIMESTAMP and the HOSTNAME. The TIMESTAMP will immediately follow the trailing ">" from the PRI part and single space characters MUST follow each of the TIMESTAMP and HOSTNAME fields. HOSTNAME will contain the hostname, as it knows itself. If it does not have a hostname, then it will contain its own IP address. If a device has multiple IP addresses, it has usually been seen to use the IP address from which the message is transmitted.

Filterlog log messages sent over syslog looks like this;
Nov 30 10:52:35 filterlog:
9,16777216,,1000000103,em0,match,block,in,4,0x0,,54,15133,0,none,6,tcp,40,x.x.x.x,x.x.x.x,48224,7547,0,S,1482288191,,37965,,
The field <Hostname> is not sent.

Some of the syslog sent from pfsense also includes that;
Dec 1 07:52:09 pfsense.effnet nginx: 10.5.10.105 - - [01/Dec/2016:07:52:09 +0100] "GET /widgets/widgets/suricata_alerts.widget.php?getNewAlerts=1480575129371 HTTP/1.1" 200 199 "https://10.5.20.1/" "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/54.0.2840.99 Safari/537.36"

If this is considered as "not a bug", the web page https://doc.pfsense.org/index.php/Filter_Log_Format_for_pfSense_2.2 should be updated accordingly and the the behavor of the other messages transmitted from pfsense which includes the hostname should be changed.

History

#1 - 07/06/2018 03:04 PM - Darren Spruell

Idar Lund wrote:

> If this is considered as "not a bug", the web page https://doc.pfsense.org/index.php/Filter_Log_Format_for_pfSense_2.2 should be updated accordingly and the the behavor of the other messages transmitted from pfsense which includes the hostname should be changed.

Could we consider a more aggressive stance on this and consider treating it as a bug? Besides the general sanity of having a syslog event state clearly the hostname that emitted it, it eases downstream complications like parsers of received log processors. Is there a great reason to require the hostname be left out?

Syslog receivers often expect syslog events to be well-formed, e.g. following the referenced RFC3164 format. For example, Logstash supports native filters for SYSLOGBASE and SYSLOGBASE2 formats, both of which require hostnames to be present in the header. The absence of this requires consumers to define custom log format filters. In fact the default grok pattern on the Logstash syslog input plugin expects this.

As Idar Lund points out there is at least one other log format in the syslog stream from pfSense that *does* include the hostname in the header: Nginx. I haven't identified any other logging that includes it, but this introduces an inconsistency in the logging (and this again puts some burden of complexity on downstream log parsers). While it may be possible to modify Nginx's logging to remove that field and unify the syslog event format, perhaps it would be better to go the other way and address this by enhancing the header on all the other logs to include the hostname.

#2 - 07/09/2018 02:34 AM - Idar Lund

I agree with Darren. This should be treated as a bug and the best solution is to add hostname to the syslog messages being sent from pfsense. The reason for adding the last line (If this is considered as "not a bug" [...]) in this bug report is because my first bug report on this got rejected before even a discussion on the topic was done: https://redmine.pfsense.org/issues/6975

#3 - 03/10/2019 05:37 PM - Daniel B

This is clearly a bug, as PfSense is not sending valid syslog messages. It also affects Graylog (3.0). We have to use a raw text input, and manually parse messages because of the missing hostname field. For graylog, if we use a syslog input, we get "filterlog" as message source instead of the host.

#4 - 03/10/2019 06:41 PM - Jim Pingle

If it's a bug, it's a bug in FreeBSD -- we use their syslogd and that's how it behaves. The default behavior is to generate rfc3164 format messages, which is what we set.

You have two potential options here:

1. Use syslog-ng, which may do what you want
2. Submit a PR to FreeBSD for them to resolve the behavior in syslogd.

#5 - 03/11/2019 03:04 PM - Daniel B

A bug is already opened upstream, see https://bugs.freebsd.org/bugzilla/show_bug.cgi?id=194231

#6 - 03/11/2019 03:19 PM - Jim Pingle

Then that is where you need to direct your attention. Comment there and let the FreeBSD developers know that it's a problem.

#7 - 08/14/2019 02:31 PM - Jim Pingle

- Category set to Logging

#8 - 07/31/2020 02:47 AM - Darren Spruell

Jim Pingle wrote:

> If it's a bug, it's a bug in FreeBSD -- we use their syslogd and that's how it behaves. The default behavior is to generate rfc3164 format messages, which is what we set.
>
> You have two potential options here:
>
> 1. Use syslog-ng, which may do what you want
> 2. Submit a PR to FreeBSD for them to resolve the behavior in syslogd.

Just looking at what I think is current status:

https://bugs.freebsd.org/bugzilla/show_bug.cgi?id=220246

This make mention of syslogd having been updated to support RFC 5424 (IETF syslog) message format. It was mentioned as being in 11-STABLE and HEAD. Looking at 11.3-RELEASE syslogd(8) (https://www.freebsd.org/cgi/man.cgi?query=syslogd&manpath=FreeBSD+11.3-RELEASE) the -O option is documented, and looking at the usage output for the syslogd binary on pfSense 2.4.5-RELEASE, it includes *-O format*.

My read is that upstream still has a buggy RFC 3164 mode in syslogd, and that's not likely to change because of backward compatibility reasons. However if the RFC 5424 option could be tested (-O rfc5424), there's a possibility it sends the HOSTNAME field to remote servers and resolves the issue. Exposing a way for users to set the desired output format in pfSense then would be really helpful.

(BTW: is there a supported way in pfSense to modify syslogd_flags to test this currently)?

#9 - 07/31/2020 06:33 AM - Jim Pingle

- Status changed from New to Feedback

An RFC 5424 option was added to 2.5.0 almost a year ago, you can test it there: #9808