

Webhooks Input - Bug #2175

Webhook doesn't start sometimes: WebServer instance has no attribute 'server'

12/20/2017 10:24 PM - Luke Murphey

Status:	Closed	Start date:	12/20/2017
Priority:	Normal	Due date:	
Assignee:	Luke Murphey	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	1.2.1		
Description			
<div>2017-12-20 06:16:59,876 ERROR Execution failed</div> <div>Traceback (most recent call last):</div> <div>File "/opt/splunk/etc/apps/webhooks_input/bin/webhooks_input_app/modular_input.py", line 1809, in execute</div> <div>self.do_run(in_stream, log_exception_and_continue=True)</div> <div>File "/opt/splunk/etc/apps/webhooks_input/bin/webhooks_input_app/modular_input.py", line 1690, in do_run</div> <div>self.run(stanza, cleaned_params, input_config)</div> <div>File "/opt/splunk/etc/apps/webhooks_input/bin/webhook.py", line 293, in run</div> <div>httpd.start_serving()</div> <div>File "/opt/splunk/etc/apps/webhooks_input/bin/webhook.py", line 193, in start_serving</div> <div>self.server.serve_forever()</div> <div>AttributeError: WebServer instance has no attribute 'server'</div>			

Associated revisions

Revision 66 - 12/21/2017 06:29 AM - lukemurphey

Added better logging to help determine why the port could not be opened

Reference #2175

Revision 68 - 12/23/2017 12:12 AM - lukemurphey

Increasing the amount of time the input will wait until giving it

Reference #2175

Revision 69 - 10/10/2018 04:21 AM - lukemurphey

Adding logging to help figure out why the input sometimes fails

Also added an if statement to prevent exception

Reference #2175

Revision 74 - 10/15/2018 09:58 PM - lukemurphey

Making input smarter such that it won't try to restart the same stanza

Reference #2175

Revision 75 - 10/16/2018 07:08 PM - lukemurphey

Adding single_instance mode support

Reference #2175

Revision 76 - 10/16/2018 08:05 PM - lukemurphey

Returning to having the input try several times to restart the server

Reference #2175

History

#1 - 12/20/2017 10:25 PM - Luke Murphey

- Target version set to 1.2.1

#2 - 12/21/2017 06:29 AM - Luke Murphey

I added better logging so I can determine why this is happening.

#3 - 12/23/2017 01:04 AM - Luke Murphey

It appears that the input is attempting to be executed twice:

```
2017-12-23 00:18:57,902 INFO Starting server on port=8081, path=u'\\ifttt\\.*', cert_file=None, key_file=None
2017-12-23 00:18:57,903 INFO The web-server could not yet be started, attempt 0 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:19:17,909 INFO The web-server could not yet be started, attempt 1 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:19:37,930 INFO The web-server could not yet be started, attempt 2 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:19:57,950 INFO The web-server could not yet be started, attempt 3 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:20:17,970 INFO The web-server could not yet be started, attempt 4 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:20:37,990 INFO The web-server could not yet be started, attempt 5 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:20:57,999 INFO The web-server could not yet be started, attempt 6 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:21:18,016 INFO The web-server could not yet be started, attempt 7 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:21:38,036 INFO The web-server could not yet be started, attempt 8 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:21:54,353 INFO Starting server on port=8081, path=u'\\ifttt\\.*', cert_file=None, key_file=None
2017-12-23 00:21:54,354 INFO The web-server could not yet be started, attempt 0 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:21:58,057 INFO The web-server could not yet be started, attempt 9 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:22:14,374 INFO The web-server could not yet be started, attempt 1 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:22:18,076 INFO The web-server could not yet be started, attempt 10 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:22:34,393 INFO The web-server could not yet be started, attempt 2 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:22:38,092 INFO The web-server could not yet be started, attempt 11 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:22:54,403 INFO The web-server could not yet be started, attempt 3 of 60, reason="[Errno 98] Address already in use"
2017-12-23 00:22:58,111 INFO The web-server could not yet be started, attempt 12 of 60, reason="[Errno 98] Address already in use"
```

#4 - 10/10/2018 04:09 AM - Luke Murphey

Observations:

- The error only appears on startup, not when restarting the input
- Extending the time for checking if the port is up doesn't help
- You cannot shut down the server by disabling the input. It keeps trying to open the port.
- The input from the previous execution appears to still be running even though Splunk was restarted
- The input appears to run under the main Splunk process on Mac. However, I see it not running under Splunk on Linux.
 - Making a new input does indeed show up under the parent process of the shell.
- I see an input try until attempt 58 and then I see a message saying that the port was successfully opened but the one opened is on a different PID.

Questions:

- Why does the error for server not being defined appear?
 - This was due to an internal problem that has been fixed.
- Is the issue what the input loops and keeps trying to run it again?
 - Appears not to be since I added checking and it still fails. Plus, it is another PID.
- Does enabling and then disabling the input fix it?
 - The input does shut down
- Does Linux see the ppid and be able to tell when the input ought to stop?
- Why does Splunk try to start it twice?
- Should I avoid the loop start entirely and use the input sleep to start the input?
 - Could inly do this in single instance mode
- Do I need the exit() after execute()?
 - It seems like it wouldn't because the process is clearly not stopping; it appears to stay running
- Why doesn't the run loop happen repeatedly?
 - This is running in multi-instance mode. Splunk manages opening the input.
 - The start server function seems to pause the code
- Would single instance mode help? Currently, it is relying on Splunk to manage this.
- Would adding a loop fix it?

2018-10-08 22:49:35,260 ERROR Execution failed

Traceback (most recent call last):

```
File "/opt/splunk/etc/apps/webhooks_input/bin/webhooks_input_app/modular_input.py", line 1809, in execute
    self.do_run(in_stream, log_exception_and_continue=True)
File "/opt/splunk/etc/apps/webhooks_input/bin/webhooks_input_app/modular_input.py", line 1690, in do_run
    self.run(stanza, cleaned_params, input_config)
File "/opt/splunk/etc/apps/webhooks_input/bin/webhook.py", line 293, in run
    httpd.start_serving()
File "/opt/splunk/etc/apps/webhooks_input/bin/webhook.py", line 193, in start_serving
    self.server.serve_forever()
```

AttributeError: WebServer instance has no attribute 'server'

Possible causes:

#5 - 10/10/2018 05:47 AM - Luke Murphey

- Status changed from New to In Progress

#6 - 10/15/2018 09:25 PM - Luke Murphey

The port is open on process 528:

tcp	0	0 0.0.0.0:8081	0.0.0.0:*	LISTEN	528/python
-----	---	----------------	-----------	--------	------------

The process looked like was disattached from Splunk.

I would have thought this was from the prior execution of the script. However, it looks like that may be a process that is left open from when the input starts when Splunk was started.

#7 - 10/15/2018 11:02 PM - Luke Murphey

See base class here; https://github.com/LukeMurphey/splunk-modular-input/blob/master/src/modular_input_base_class.py

#9 - 10/16/2018 08:20 PM - Luke Murphey

- Status changed from In Progress to Closed