Network Tools - Bug #2348

Lookups do not work

12/12/2018 07:04 PM - Luke Murphey

Status:	Closed	**Start date:**	12/12/2018
Priority:	Normal	**Due date:**	
Assignee:	Luke Murphey	**% Done:**	100%
Category:		**Estimated time:**	0.00 hour
Target version:	1.2.9		

Description

Associated revisions

Revision 281 - 12/18/2018 08:45 PM - lukemurphey

Improving likelihood that lookups work

Reference #2348

Revision 282 - 12/18/2018 10:01 PM - lukemurphey

Making lookups work with 7.2.1

I no longer add fields to the header since this apparently causes Splunk
to not match up the output with the input

Reference #2348

Revision 283 - 12/18/2018 10:04 PM - lukemurphey

Adding comments and disabling to writing up existing fields

Reference #2348

History

#1 - 12/12/2018 07:34 PM - Luke Murphey

Observations:

- Splunk sees the field names but not the values of output from a scripted lookup
- It does show values for new columns in the case of the external lookup
- A field in the lookup that is rewritten by the script is viewed as empty in the UI
  - However, the external_lookup example does this correctly
- The rows are being written
- Output is being received
- This works on 7.1.4 and earlier per the customer but I found that it doesn't work on 7.1.2
- I replaced ping lookup with the external lookup example, it then outputs the columns
  - | inputlookup append=t domains.csv | lookup ping clienthost AS host
  - | inputlookup append=t domains.csv | lookup dnslookup clienthost AS host
- Splunk seems to not output fields unless the fields are in the list of arguments it accepts
- The example does output a field that doesn't appear in the lookup contents
  - | inputlookup append=t testfile.csv | lookup dnslookup clienthost as host
- Splunk swaps out the arguments of the
- The NT search command doesn't seem to look at the incoming args for the header
- If I include the field in the search, it gets listed (though using the value from the lookup file)

    - Includes test field: | inputlookup append=t testfile.csv | lookup test_lookup clienthost as host test AS test
    - Does not include test field: | inputlookup append=t testfile.csv | lookup test_lookup clienthost as host
- I am able to add fields to the incoming input in an example
- Removing the row-writes changes nothing; it is as if Splunk is ignoring the output entirely
- The following return different results:
    - | inputlookup append=t testfile.csv | lookup ping host as test2
    - | inputlookup append=t testfile.csv | lookup ping host as test

Questions:

- Where is an example of scripted lookups?
    - http://docs.splunk.com/Documentation/Splunk/7.2.1/Knowledge/Configureexternallookups
- Is the raw output the problem?
    - Outputting easily parsable content still doesn't work
- Why does the built-in one work? Could be because it uses the same header and just adds fields to it?
- Is this a platform issue?
    - It doesn't work on 7.2.0 on Unix
- Does setting the type help?
    - executable and python doesn't help
- Do the examples work?
    - https://answers.splunk.com/answers/145561/how-to-script-a-lookup-in-python.html
    - External works: | inputlookup append=t hosts.csv | lookup dnslookup clienthost AS host
- Does Splunk change the data or the arguments when "AS" is used?
    - Splunk seems to swap out the argument in the call
- What happens if I have the sample use the same header that was provided but add one field?
- Why does the external lookup example allow a new field to be added? Is it because it is in the original lookup or because it is a field name that is included in the command-line?
    - This works: | inputlookup append=t ping_hosts.csv | lookup test_lookup clienthost as host
- What happens if I just have the ping command output the original fields back?

#2 - 12/12/2018 07:41 PM - Luke Murphey

Simple script to test output of command:

```
export SPLUNK_HOME=/Users/lmurphey/Splunk/721
export PYTHONPATH=$SPLUNK_HOME/lib/python2.7
export SPLUNK_DB=$SPLUNK_HOME/var/lib
export SPLUNK_ETC=$SPLUNK_HOME/etc

cat $SPLUNK_HOME/etc/apps/network_tools/lookups/hosts.csv | $SPLUNK_HOME/bin/python $SPLUNK_HOME/etc/apps/netw
ork_tools/bin/ping_lookup.py host
```

#3 - 12/12/2018 07:53 PM - Luke Murphey

```
cat $SPLUNK_HOME/etc/apps/network_tools/lookups/hosts.csv | $SPLUNK_HOME/bin/python $SPLUNK_HOME/etc/apps/netw
ork_tools/bin/ping_lookup.py host

received,jitter,packet_loss,min_ping,avg_ping,return_code,host,max_ping,raw_output,sent
1,0.000,0.0,1.148,1.148,0,10.0.0.6,1.148,asd,1
1,0.000,0.0,1.148,1.148,0,10.0.0.6,1.148,asd,1
```

#4 - 12/12/2018 07:56 PM - Luke Murphey

```
cat $SPLUNK_HOME/etc/apps/network_tools/lookups/test.csv | $SPLUNK_HOME/bin/python $SPLUNK_HOME/etc/system/bin
/external_lookup.py host ip

host,ip
work.com,127.0.0.1
```

#5 - 12/18/2018 01:05 AM - Luke Murphey

This outputs the "test" field with the output field set to "AAA":

```
import csv
import sys

def doit():
    infile = sys.stdin
    outfile = sys.stdout

    r = csv.DictReader(infile)

    header = r.fieldnames
    header.append('test')

    w = csv.DictWriter(outfile, fieldnames=header)
    w.writeheader()

    for result in r:
        result['test'] = 'AAA'

        w.writerow(result)

doit()
```

This does not:

```
import csv
import sys

def doit():
    infile = sys.stdin
    outfile = sys.stdout

    r = csv.DictReader(infile)
```

```
    header = r.fieldnames
    header.append('test')

    w = csv.DictWriter(outfile, fieldnames=header)
    w.writeheader()

    for result in r:
        for key, value in result.items():
            result[key] = 'BBB'

        result['test'] = 'AAA'

        w.writerow(result)

doit()
```

#6 - 12/18/2018 01:10 AM - Luke Murphey

In the above code:

These includes the "test" field:

```
    for result in r:
        for key, value in result.items():
            result[key] = value

        result['test'] = 'AAA'
```

```
    for result in r:
        result['test'] = 'AAA'
```

This doesn't:

```
    for result in r:
        for key, value in result.items():
            result[key] = 'AAA'

        result['test'] = 'BBB'
```

```
    for result in r:
        for key, value in result.items():
            result[key] = str(value) + 'A'

        result['test'] = 'BBB'
```

#7 - 12/18/2018 07:19 PM - Luke Murphey

I found that Splunk was aggregating my results as mkv fields under the host "10.0.0.6" even though I didn't output the values as MKV fields. It turns out that Splunk is matching up the imput field (host) with the output field host. When it sees multiple rows with the same host field, it matches up the rows and produces an MKV with all of the values.

Thus, this worked:

```
        output = {
            'raw_output': 'NOW!',
            'min_ping': '1.148',
            'received': '1',
            'jitter': '0.000',
            'packet_loss': '0.0',
            'sent': '1',
            'max_ping': '1.148',
            'avg_ping': '1.148',
            'return_code': "0"
        }
```

```
        for k in output:
            if k not in header:
                header.append(k)
```

This produced MKV fields since I had host hardcoded:

```
        output = {
            'raw_output': 'NOW!',
            'min_ping': '1.148',
            'host': '10.0.0.6',
            'received': '1',
            'jitter': '0.000',
            'packet_loss': '0.0',
            'sent': '1',
            'max_ping': '1.148',
            'avg_ping': '1.148',
            'return_code': "0"
        }
```

```
        for k in output:
            if k not in header:
                header.append(k)
```

#8 - 12/18/2018 08:41 PM - Luke Murphey

I added an argument for for not adding fields if they exist with an empty value. This worked on 7.1.2 but not on 7.2.1.

#9 - 12/18/2018 10:23 PM - Luke Murphey

- Status changed from New to Closed

- % Done changed from 0 to 100