

## GlideinWMS - Bug #2984

Support # 2139 (Closed): Check support for platform SL6

Bug # 2634 (Closed): Migrate popen calls to subprocess for SL6 compatibility

### Problematic condorExe behavior

10/04/2012 03:56 PM - Igor Sfiligoi

<b>Status:</b>	Closed	<b>Start date:</b>	10/04/2012
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Burt Holzman	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v2_7	<b>Spent time:</b>	0.00 hour
<b>First Occurred:</b>		<b>Stakeholders:</b>	
<b>Occurs In:</b>			
<b>Description</b>			
<p>The OSG gfactory team has noticed a weird behavior for the python processes when executing external commands. The cuprit seem to be the condorExec.iexe_cmd implementation (v2plus-related flavor).</p> <p>The major lead was the fact that the python process would fork itself, run for a short period of time, and then replace itself with the executed command.</p> <p>FYI: This seems to be using a significant amount of CPU on a busy factory.</p> <p>However, the real problem is that if you try to send a sigterm to this process before it replaces itself with the command to execute, the process flow gets completely screwed up!</p> <p>The child will think it is the parent, and will enter the main factory loop!</p>			

### History

#### #1 - 10/04/2012 03:59 PM - Igor Sfiligoi

Here is a snippet of an offending process.

PID 25226 is the parent, and PID 18546 is the child who was sent the SIGTERM.

```
[2012-10-04T19:33:50+00:00 25226] Iteration at Thu Oct 4 19:33:50 2012
[2012-10-04T19:40:14+00:00 25226] Found 1 tasks to work on using existing factory key.
[2012-10-04T19:40:14+00:00 25226] Found 1 total tasks to work on
[2012-10-04T19:40:16+00:00 18546] Schedd schedd_glideins5@glidein-itb.grid.iu.edu not responding, skipping
[2012-10-04T19:40:16+00:00 18546] WARNING: getCondorQData failed: Unexpected Error running '/usr/local/glidecondor/bin/condor_q -format "%i" "JobStatus" -format "%s" "GridJobStatus" -format "%i" "ServerTime" -format "%i" "EnteredCurrentStatus" -format "%s" "GlideinX509Identifier" -format "%i" "HoldReasonCode" -format "%i" "HoldReasonSubCode" -format "%i" "NumSystemHolds" -format "%s" "GlideinFrontendName" -format "%s" "GlideinClient" -format "%s" "GlideinX509SecurityClass" -format "%i" "ClusterId" -format "%i" "ProcId" -xml -constraint '(GlideinFactory =?= "ITBGOC") && (GlideinName =?= "v1_1") && (GlideinEntryName =?= "CDF_IT_CNAF_ce01") && (GlideinX509Identifier != UNDEFINED)''
Stdout:None
Stderr:None
Exception OSError: Received signal 15
[2012-10-04T19:40:16+00:00 18546] Advertize
[2012-10-04T19:40:18+00:00 18546] Writing stats
[2012-10-04T19:40:18+00:00 18546] log_stats written
[2012-10-04T19:40:18+00:00 18546] qc_stats written
[2012-10-04T19:40:19+00:00 18546] rrd_stats written
[2012-10-04T19:40:19+00:00 18546] Sleep 60s
[2012-10-04T19:40:41+00:00 18546] Received signal...exit
[2012-10-04T19:40:41+00:00 18546] Dying
[2012-10-04T19:40:41+00:00 25226] Schedd schedd_glideins5@glidein-itb.grid.iu.edu not responding, skipping
[2012-10-04T19:40:41+00:00 25226] WARNING: getCondorQData failed: Error running '/usr/local/glidecondor/bin/condor_q -format "%i" "JobStatus" -format "%s" "GridJobStatus" -format "%i" "ServerTime" -format "%i" "EnteredCurrentStatus" -format "%s" "GlideinX509Identifier" -format "%i" "HoldReasonCode" -format "%i" "HoldReasonSubCode" -format "%i" "NumSystemHolds" -format "%s" "GlideinFrontendName" -format "%s" "GlideinClient" -format "%s" "GlideinX509SecurityClass" -format "%i" "ClusterId" -format "%i" "ProcId" -xml -constraint '(GlideinFactory =?= "ITBGOC") && (GlideinName =?= "v1_1") && (GlideinEntryName =?= "CDF_IT_CNAF_ce01") && (GlideinX509Identifier != UNDEFINED)''
```

```
code 1:Traceback (most recent call last):
  File "/home/gfactory/glideinWMS/factory/glideFactoryEntry.py", line 1026, in ?
    main(sys.argv[1],int(sys.argv[2]),int(sys.argv[3]),sys.argv[4],sys.argv[5])
  File "/home/gfactory/glideinWMS/factory/glideFactoryEntry.py", line 1011, in main
    pid_obj.relinquish()
  File "/home/gfactory/glideinWMS/factory/./lib/pidSupport.py", line 81, in relinquish
    self.fd.seek(0)
IOError: [Errno 9] Bad file descriptor
close failed: [Errno 9] Bad file descriptor
```

```
[2012-10-04T19:40:41+00:00 25226] Advertize
[2012-10-04T19:40:42+00:00 25226] Writing stats
[2012-10-04T19:40:42+00:00 25226] log_stats written
[2012-10-04T19:40:42+00:00 25226] qc_stats written
[2012-10-04T19:40:43+00:00 25226] rrd_stats written
```

## #2 - 10/04/2012 04:03 PM - Igor Sfiligoi

It can get even worse... in this case, the child took over the entry.... the parent never unblocked (but it is still alive!)

```
[2012-10-04T19:54:27+00:00 24921] Iteration at Thu Oct 4 19:54:27 2012
[2012-10-04T19:57:11+00:00 24921] Found 0 tasks to work on using existing factory key.
[2012-10-04T19:57:11+00:00 24921] Old factory key is still valid. Trying to find work using old factory key.
[2012-10-04T19:57:15+00:00 24921] Found 3 tasks to work on using old factory key
[2012-10-04T19:57:15+00:00 24921] Found 3 total tasks to work on
[2012-10-04T19:57:16+00:00 24921] Checking downtime for frontend UCSDOSG-itb security class: frontend (entry C
MS_T3_US_Colorado_hep01).
[2012-10-04T19:57:18+00:00 26767] WARNING: Failed to update x509_proxy_0 using username feucsditb for client UC
SDOSG-itb-v1_1.cms, skipping request
[2012-10-04T19:57:18+00:00 26767] WARNING: No good proxies for UCSDOSG-itb-v1_1.cms, skipping request
[2012-10-04T19:57:18+00:00 26767] Checking downtime for frontend UCSDCMS security class: cmstpilot (entry CMS_T
3_US_Colorado_hep01).
[2012-10-04T19:57:20+00:00 26767] Client UCSD-i5_2a.itb (secid: UCSDCMS_cmstpilot) requesting 0 glideins, max r
unning 0, remove excess 'ALL'
[2012-10-04T19:57:20+00:00 26767] Params: {u'GLIDEIN_Glexec_Use': u'OPTIONAL', u'USE_MATCH_AUTH': u'True', '
GLIDEIN_Job_Max_Time': 34800, u'GLIDEIN_Collector': u'glidein-collector.t2.ucsd.edu:9620-9919'}
[2012-10-04T19:57:20+00:00 26767] Decrypted Param Names: [u'SecurityName', u'nr_x509_proxies', u'x509_proxy_
0', u'x509_proxy_0_identififier', u'x509_proxy_0_security_class']
...
[2012-10-04T21:01:23+00:00 26767] Sleep 60s
[2012-10-04T21:02:23+00:00 26767] Iteration at Thu Oct 4 21:02:23 2012
```

```
2012-10-04 21:02:04 UTC [gfactory@glidein-itb:~]$ ps -ef |grep 24921
gfactory 24921 24526 0 19:53 pts/1 00:00:01 /usr/bin/python /home/gfactory/glideinWMS/factory/glideFactory
Entry.py 24526 60 5 /usr/local/share/gfactory/glideinsubmit/glidein_v1_1 CMS_T3_US_Colorado_hep01
gfactory 26767 24921 0 19:57 pts/1 00:00:15 /usr/bin/python /home/gfactory/glideinWMS/factory/glideFactory
Entry.py 24526 60 5 /usr/local/share/gfactory/glideinsubmit/glidein_v1_1 CMS_T3_US_Colorado_hep01
```

## #3 - 10/04/2012 04:08 PM - Igor Sfiligoi

- File a.py added

- File condorExe.py added

The behavior is not too difficult to reproduce.

I created a simple test script, a.py, attached, and instrumented condorExe to print out the pids.

As can be seen, both processes print something out when a SIGTERM is sent:

```
2012-10-04 20:47:57 UTC [gfactory@glidein-itb:~/glideinWMS/lib]$ ./a.py
22425 start1
22425 about to popen /usr/local/glidecondor/bin/./././././bin/sleep 3
22425 done popen /usr/local/glidecondor/bin/./././././bin/sleep 3
22425 before while
22425 after while
22425 about to exit
Traceback (most recent call last):
  File "./a.py", line 13, in ?
    a=condorExe.exe_cmd("./././././bin/sleep","3")
  File "/home/gfactory/glideinWMS/lib/condorExe.py", line 59, in exe_cmd
    return iexe_cmd(cmd,stdin_data,env)
  File "/home/gfactory/glideinWMS/lib/condorExe.py", line 170, in iexe_cmd
    raise ExeError, "%i Error running '%s'\n\ncode %i:%s" % (os.getpid(),cmd, os.WEXITSTATUS(exitStatus), ".joi
```

```

n(error_lines))
condorExe.ExeError: 22425 Error running '/usr/local glidecondor/bin/../../../../bin/sleep 3'
code 1:Traceback (most recent call last):
  File "./a.py", line 13, in ?
    a=condorExe.exe_cmd("../../../../bin/sleep","3")
  File "/home/gfactory/glideinWMS/lib/condorExe.py", line 59, in exe_cmd
    return iexe_cmd(cmd,stdin_data,env)
  File "/home/gfactory/glideinWMS/lib/condorExe.py", line 165, in iexe_cmd
    raise ExeError, "%i Unexpected Error running '%s'\nStdout:%s\nStderr:%s\n" \
condorExe.ExeError: 22427 Unexpected Error running '/usr/local glidecondor/bin/../../../../bin/sleep 3'
Stdout:None
Stderr:None
Exception OSError: Received signal 15

2012-10-04 20:43:55 UTC [gfactory@glidein-itb:~]$ ps -ef |grep a.py |grep -v 499 |head -1 |tee a.txt |awk '{pr
int $2}' | xargs kill -TERM

```

**#4 - 10/05/2012 12:59 PM - Burt Holzman**

- Status changed from New to Assigned
- Target version set to v2\_7\_x
- Parent task set to #2634

I dug deeper here. The main issue is that popen always invokes the shell (/bin/sh -c cmd). This shouldn't be necessary. Additionally, i18n can make execing bash slow, too. As a workaround, Igor's trying to export LANG=C on the factories and reports that things look better.

The real answer is to review 2634 fully and merge it. I'll leave this open until that happens.

**#5 - 11/19/2012 02:54 PM - Burt Holzman**

- Status changed from Assigned to Closed

**#6 - 11/19/2012 02:54 PM - Burt Holzman**

- Status changed from Closed to Resolved

**#7 - 11/19/2012 03:04 PM - Parag Mhashilkar**

- Target version changed from v2\_7\_x to 293

[#2634](#) fixes this issue

**#8 - 03/18/2013 01:20 PM - Parag Mhashilkar**

- Target version changed from 293 to v2\_7

**#9 - 03/18/2013 01:56 PM - Parag Mhashilkar**

- Status changed from Resolved to Closed

**Files**

a.py	352 Bytes	10/04/2012	Igor Sfiligoi
condorExe.py	8.17 KB	10/04/2012	Igor Sfiligoi