

GlideinWMS - Bug #19524

Protect glidien against duplicate values in Factory and Frontend

03/28/2018 04:33 PM - Marco Mambelli

Status:	New	Start date:	03/28/2018
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
First Occurred:		Stakeholders:	
Occurs In:	v3_2_23		

Description

The glidiens use a file called glidein_config to receive configuration information and to communicate among the different shell scripts. This file is expected to have lines with key, value pairs separated by one or more spaces. All keys are expected to be unique.

A misconfiguration highlighted a problem:

- values are selected w/ lines like: `error_gen=`grep '^ERROR_GEN_PATH ' $config_file | awk '{print $2}'``
- if there are duplicates this causes the result to be a string with values separated by newlines (one line per occurrence). This will cause errors in the scripts (e.g. when numbers are expected)

Here you can see the problem:

```
# --- Last Script values ---
```

```
-----  
max wall time, 7200  
84600  
/var/lib/condor/execute/dir_4841/glide_UyWdlz/main/condor_startup.sh: line 483: [: 7200  
84600: integer expression expected  
job max time, 34800  
/var/lib/condor/execute/dir_4841/glide_UyWdlz/main/condor_startup.sh: line 492: let: die_time=7200  
84600 - 370 - 120: syntax error in expression (error token is "84600 - 370 - 120")  
calculated retire time, -27600  
used param retire spread, 1200  
Retire time after spread too low (-28428), remove spread
```

This happened because both the factory and Frontend had a configuration line like:

```
<attr name="GLIDEIN_Max_Walltime" comment="opportunistic vos only get 24h" const="True" glidein_publish="True"  
job_publish="False" parameter="True" publish="True" type="int" value="84600"/>  
and
```

```
<attr name="GLIDEIN_MAX_WALLTIME" glidein_publish="False" job_publish="False" parameter="True" type="int" value="7200"/>
```

Frontends should use GLIDEIN_Job_Max_Time, so this was solved as a configuration error.

Anyway the mechanism should be more robust:

- isn't there anything to protect vs duplicate lines (if a value is set in both, should one override the other is a known way)?
- if there is, why did not work here?
- should the parse line include a protection (e.g. `grep '^ERROR_GEN_PATH ' $config_file | head -n1 | awk '{print $2}' - first or last (tail) value`)?

This ticket is: to document glidien_config mechanics, to understand what is in place to protect against duplicates, to understand why it did not work, to survey if there are values relying on multiple values of the key, evaluate if it is worth to put in place a generic mechanism or if it is better to protect at least some key variables from duplicates (or do nothing and leave as it is).