

GlideinWMS - Bug #3103

Overly complicated secondary schedd config

11/01/2012 09:35 PM - Igor Sfiligoi

Status: New	Start date: 11/01/2012
Priority: Normal	Due date:
Assignee:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: v3_x	Spent time: 0.00 hour
First Occurred:	Stakeholders:
Occurs In:	
Description Looking at the secondary schedd config created by both the ini installer and the RPMs, it seems overly complicated. We re-define the whole schedd world. Why not use the environment, instead? Like we do with the secondary collectors? We know the needed values, since we used them in the Q&A mode.	
Related issues: Related to GlideinWMS - Feature #3101: Dynamically creating the secondary sch... Closed 11/01/2012	

History

#1 - 11/01/2012 09:40 PM - Igor Sfiligoi

We can replace

```
SCHEDDGLIDEINS5_ARGS = -local-name scheddglideins5
SCHEDD.SCHEDDGLIDEINS5.SCHEDD_NAME = schedd_glideins5
SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR = $(LOCAL_DIR)/$(SCHEDD.SCHEDDGLIDEINS5.SCHEDD_NAME)
SCHEDD.SCHEDDGLIDEINS5.EXECUTE = $(SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR)/execute
SCHEDD.SCHEDDGLIDEINS5.LOCK = $(SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR)/lock
SCHEDD.SCHEDDGLIDEINS5.LOG = $(SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR)/log
SCHEDD.SCHEDDGLIDEINS5.PROCD_ADDRESS = $(SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR)/procd_pipe
SCHEDD.SCHEDDGLIDEINS5.SPOOL = $(SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR)/spool
SCHEDD.SCHEDDGLIDEINS5.JOB_QUEUE_LOG = $(SCHEDD.SCHEDDGLIDEINS5.SPOOL)/job_queue.log
SCHEDD.SCHEDDGLIDEINS5.SCHEDD_ADDRESS_FILE = $(SCHEDD.SCHEDDGLIDEINS5.SPOOL)/.schedd_address
SCHEDD.SCHEDDGLIDEINS5.SCHEDD_DAEMON_AD_FILE = $(SCHEDD.SCHEDDGLIDEINS5.SPOOL)/.schedd_classad
SCHEDDGLIDEINS5_LOCAL_DIR_STRING = "$(SCHEDD.SCHEDDGLIDEINS5.LOCAL_DIR)"
SCHEDD.SCHEDDGLIDEINS5.SCHEDD_EXPRS = LOCAL_DIR_STRING
```

```
SCHEDDGLIDEINS5_ENVIRONMENT = "_CONDOR_GRIDMANAGER_LOG=$(LOG)/GridManagerLog,$(SCHEDD.SCHEDDGLIDEINS5.SCHEDD_NAME) .$(USERNAME)"
```

with just

```
SCHEDDGLIDEINS5_ENVIRONMENT = "_CONDOR_SCHEDD_NAME=schedd_glideins5 _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glideins5 _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glideins5/lock"
```

(copied pretty much verbatim from what we had in the scripts created by Q&A)

#2 - 11/02/2012 08:53 AM - Burt Holzman

This is what John had to do to get it to work -- if your prescription works then we should use it (as long as it is wholly handled by the config file and condor manages the schedd startup, unlike the Q&A which launched them one at a time in a for loop)

#3 - 11/05/2012 08:47 AM - John Weigand

Igor,

I assume you have this working in a test environment. Can you be more specific as to exactly what the complete set of attributes should be for a second schedd. I tried and failed based on what you had in this ticket. Likely me missing something. I would really like this to be simpler.

Thanks
John Weigand

#4 - 11/05/2012 01:30 PM - Igor Sfiligoi

Ops... forgot about the shared daemon fix.

Looking at the latest Q&A installer, the right string to add should be

```
SCHEDDGLIDEINS5_ENVIRONMENT = \  
"_CONDOR_SCHEDD_NAME=schedd_glideins5 _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glideins5 _CONDOR_LOCK=$(LOCAL_DIR)  
)/schedd_glideins5/lock _CONDOR_SHARED_PORT_DAEMON_AD_FILE=$(LOG)/shared_port_ad _CONDOR_DAEMON_SOCKET_DIR=$(L  
OG)/daemon_sock"
```

Sorry about that,
Igor

#5 - 11/06/2012 08:55 AM - Parag Mhashilkar

I really don't see how one solution is more complicated than other. One defines a variable in each line while other defines all the variables in one line and that is converted to an extremely long environment line.

#6 - 11/06/2012 09:26 AM - Igor Sfiligoi

The difference is WHAT we set;

- in the "Q&A mode", we set the base directories for the SCHEDD (+the preserve the shared port settings).
- in "current RPM mode". we set ALL the variables the SCHEDD may need to use (on the plus side, no need to tamper with shared port)

It is just much harder to maintain the "ALL options" variant;
we were hit already once, when Condor added an additional parameter to the schedd, and we did not define it... would not happen in the "Q&A mode".

Igor

#7 - 11/06/2012 09:45 AM - Parag Mhashilkar

Igor Sfiligoi wrote:

The difference is WHAT we set;

- in the "Q&A mode", we set the base directories for the SCHEDD (+the preserve the shared port settings).
- in "current RPM mode". we set ALL the variables the SCHEDD may need to use (on the plus side, no need to tamper with shared port)

This is not about Q&A v/s RPM v/s ini installer issue. This is about how we should configure secondary schedd/collector and once we decide which way we want to go, installer(s) should be changed accordingly.

It is just much harder to maintain the "ALL options" variant;
we were hit already once, when Condor added an additional parameter to the schedd, and we did not define it... would not happen in the "Q&A mode".

If you are talking about lossing queue log, that was a nasty bug in condor and they fixed it. It's just a coincidence that in Q&A mode you didn't see it.

#8 - 11/06/2012 09:46 AM - Igor Sfiligoi

BTW: I think I have found a way to simplify it further (and partially address Parag's concern).

We can define the shared port daemon settings only once, and then reuse them for all the schedds.
Unfortunately, concatenating strings does not work in the condor config, so we need to do a 2 step process.

```
SSP_SHPORT= _CONDOR_SHARED_PORT_DAEMON_AD_FILE=$(SHARED_PORT_DAEMON_AD_FILE) _CONDOR_DAEMON_SOCKET_DIR=$(DAEMON  
N_SOCKET_DIR)  
...  
SCHEDDGLIDEIN1_ES = _CONDOR_SCHEDD_NAME=schedd_glidein1 _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein1 _CONDOR  
_LOCK=$(LOCAL_DIR)/schedd_glidein1/lock $(SSP_SHPORT)  
SCHEDDGLIDEIN1_ENVIRONMENT = "$ (STR_SCHEDDGLIDEIN1_ES) "  
...  
SCHEDDGLIDEIN5_ES = _CONDOR_SCHEDD_NAME=schedd_glidein5 _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein5 _CONDOR  
_LOCK=$(LOCAL_DIR)/schedd_glidein5/lock $(SSP_SHPORT)  
SCHEDDGLIDEIN5_ENVIRONMENT = "$ (STR_SCHEDDGLIDEIN5_ES) "  
...
```

Now we really just define the straight minimum for each schedd ;)

#9 - 11/06/2012 09:48 AM - Igor Sfiligoi

As for "Q&A vs RPM", I agree we should sync them.

Was just using those names to make it easier to refer to one vs the other way of writing configs.

#10 - 11/06/2012 09:57 AM - Igor Sfiligoi

Sigh... the coffee is just starting to work.

Of course there is a way to inject a string into another string!

(but note that the SHPORT "string" is not really a string)

```
SSP_SHPORT= _CONDOR_SHARED_PORT_DAEMON_AD_FILE=$(SHARED_PORT_DAEMON_AD_FILE) _CONDOR_DAEMON_SOCKET_DIR=$(DAEMON_SOCKET_DIR)
...
SCHEDDGLIDEIN5_ENVIRONMENT = "_CONDOR_SCHEDD_NAME=schedd_glidein5 _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein5 _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein5/lock $(SSP_SHPORT) "
...
```

#11 - 11/06/2012 10:18 AM - Parag Mhashilkar

Just a data point after talking to Dan,

- Environment is inherited by children, whereas the other way is not.
- If you would ever want to reconfig any of these settings without restarting the daemon, this does not work the environment way.

I am perfectly fine either ways. I don't think it's going to affect the complexity by a big margin, its just necessary evil which we will have to deal with.

#12 - 11/06/2012 10:22 AM - John Weigand

Igor,

Can you you expand on that last comment about your discussion with Dan about the environment?

Mainly the 2nd bullet. I have always had a hard time understanding when a condor_reconfig works and does not work.

Keep it simple especially for me.

Thanks

John Weigand

#13 - 11/06/2012 10:26 AM - Igor Sfiligoi

Not much to say on the 2nd point from my part;

Dan just mentioned to me the same thing.

However, given that we only define directory/file locations via config/env, I don not expect the schedd would be happy we changed those without a restart anyhow.

PS: Point (1) is important though... I think the "current RPM/ini mode" does not properly configure the shadows... which the "Q&A mode i.e. env mode" does, due to child inheritance.

#14 - 11/07/2012 10:17 AM - John Weigand

Igor,

Can you put in this ticket what you consider the "final" configuration entries of a single secondary schedd?

John Weigand

#15 - 11/07/2012 01:09 PM - Igor Sfiligoi

Here is what I would put for 2 schedds, if we keep separate log directories (something I personally like a lot):

```
PRESERVE_SHPORT_EXPENV= _CONDOR_SHARED_PORT_DAEMON_AD_FILE=$(SHARED_PORT_DAEMON_AD_FILE) \
                        _CONDOR_DAEMON_SOCKET_DIR=$(DAEMON_SOCKET_DIR)

# Secondary schedd schedd_glidein1
SCHEDDGLIDEIN1 = $(SCHEDD)
SCHEDDGLIDEIN1_EXPENV = _CONDOR_SCHEDD_NAME=schedd_glidein1 \
                        _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein1 \
                        _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein1/lock
SCHEDDGLIDEIN1_ENVIRONMENT = "$ (SCHEDDGLIDEIN1_EXPENV) $ (PRESERVE_SHPORT_EXPENV) "
```

```

SEC_SCHEDD_LIST = $(SEC_SCHEDD_LIST) SCHEDDGLIDEIN1

# Secondary schedd schedd_glidein2
SCHEDDGLIDEIN2 = $(SCHEDD)
SCHEDDGLIDEIN2_EXPENV = _CONDOR_SCHEDD_NAME=schedd_glidein2 \
                        _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein2 \
                        _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein2/lock
SCHEDDGLIDEIN2_ENVIRONMENT = "$ (SCHEDDGLIDEIN2_EXPENV) $(PRESERVE_SHPORT_EXPENV) "
SEC_SCHEDD_LIST = $(SEC_SCHEDD_LIST) SCHEDDGLIDEIN2

DAEMON_LIST = $(DAEMON_LIST) $(SEC_SCHEDD_LIST)
DC_DAEMON_LIST = + $(SEC_SCHEDD_LIST)

```

If instead you want all logs in the same dir, it becomes:

```

PRESERVE_SHPORT_EXPENV= _CONDOR_SHARED_PORT_DAEMON_AD_FILE=$(SHARED_PORT_DAEMON_AD_FILE) \
                        _CONDOR_DAEMON_SOCKET_DIR=$(DAEMON_SOCKET_DIR)

# Secondary schedd schedd_glidein1
SCHEDDGLIDEIN1 = $(SCHEDD)
SCHEDDGLIDEIN1_EXPENV = _CONDOR_SCHEDD_NAME=schedd_glidein1 \
                        _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein1 \
                        _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein1/lock \
                        _CONDOR_SCHEDD_LOG=$(LOG)/SchedLog.schedd_glidein1 \
                        _CONDOR_SHADOW_LOG=$(LOG)/ShadowLog.schedd_glidein1
SCHEDDGLIDEIN1_ENVIRONMENT = "$ (SCHEDDGLIDEIN1_EXPENV) $(PRESERVE_SHPORT_EXPENV) "
SEC_SCHEDD_LIST = $(SEC_SCHEDD_LIST) SCHEDDGLIDEIN1

# Secondary schedd schedd_glidein2
SCHEDDGLIDEIN2 = $(SCHEDD)
SCHEDDGLIDEIN2_EXPENV = _CONDOR_SCHEDD_NAME=schedd_glidein2 \
                        _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein2 \
                        _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein2/lock \
                        _CONDOR_SCHEDD_LOG=$(LOG)/SchedLog.schedd_glidein2 \
                        _CONDOR_SHADOW_LOG=$(LOG)/ShadowLog.schedd_glidein2
SCHEDDGLIDEIN2_ENVIRONMENT = "$ (SCHEDDGLIDEIN2_EXPENV) $(PRESERVE_SHPORT_EXPENV) "
SEC_SCHEDD_LIST = $(SEC_SCHEDD_LIST) SCHEDDGLIDEIN2

DAEMON_LIST = $(DAEMON_LIST) $(SEC_SCHEDD_LIST)
DC_DAEMON_LIST = + $(SEC_SCHEDD_LIST)

```

PS2: I expanded the list on multiple lines, since there were some comments against long lines.

PS2: This was obtained by running the script from [#3101](#).

#16 - 11/07/2012 01:19 PM - Igor Sfiligoi

Just thought of a slight improvement for the "all logs in the same dir" scenario:

```

PRESERVE_SHPORT_EXPENV= _CONDOR_SHARED_PORT_DAEMON_AD_FILE=$(SHARED_PORT_DAEMON_AD_FILE) \
                        _CONDOR_DAEMON_SOCKET_DIR=$(DAEMON_SOCKET_DIR)

# Secondary schedd schedd_glidein1
SCHEDDGLIDEIN1 = $(SCHEDD)
SCHEDDGLIDEIN1_EXPENV = _CONDOR_SCHEDD_NAME=schedd_glidein1 \
                        _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein1 \
                        _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein1/lock \
                        _CONDOR_SCHEDD_LOG=$(SCHEDD_LOG).schedd_glidein1 \
                        _CONDOR_SHADOW_LOG=$(SHADOW_LOG).schedd_glidein1
SCHEDDGLIDEIN1_ENVIRONMENT = "$ (SCHEDDGLIDEIN1_EXPENV) $(PRESERVE_SHPORT_EXPENV) "
SEC_SCHEDD_LIST = $(SEC_SCHEDD_LIST) SCHEDDGLIDEIN1

# Secondary schedd schedd_glidein2
SCHEDDGLIDEIN2 = $(SCHEDD)
SCHEDDGLIDEIN2_EXPENV = _CONDOR_SCHEDD_NAME=schedd_glidein2 \
                        _CONDOR_LOCAL_DIR=$(LOCAL_DIR)/schedd_glidein2 \
                        _CONDOR_LOCK=$(LOCAL_DIR)/schedd_glidein2/lock \
                        _CONDOR_SCHEDD_LOG=$(SCHEDD_LOG).schedd_glidein2 \
                        _CONDOR_SHADOW_LOG=$(SHADOW_LOG).schedd_glidein2
SCHEDDGLIDEIN2_ENVIRONMENT = "$ (SCHEDDGLIDEIN2_EXPENV) $(PRESERVE_SHPORT_EXPENV) "
SEC_SCHEDD_LIST = $(SEC_SCHEDD_LIST) SCHEDDGLIDEIN2

DAEMON_LIST = $(DAEMON_LIST) $(SEC_SCHEDD_LIST)
DC_DAEMON_LIST = + $(SEC_SCHEDD_LIST)

```

#17 - 01/07/2013 02:44 PM - Burt Holzman

- *Target version set to v3_1*

#18 - 04/30/2013 08:48 AM - Burt Holzman

- *Assignee changed from Igor Sfiligoi to John Weigand*

- *Target version changed from v3_1 to v3_x*

Since we won't ship the secondary schedds by default (but rely on the tool to build it), we just need to confirm that the tool builds the simpler config for 3.x.

#19 - 05/19/2015 01:53 PM - Parag Mhashilkar

- *Assignee deleted (John Weigand)*