

GlideinWMS - Feature #2172

Making config file location optional during reconfig

11/17/2011 11:34 AM - Igor Sfiligoi

Status: Closed	Start date: 01/23/2012
Priority: Low	Due date:
Assignee: Krista Larson	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: v2_6	Spent time: 0.50 hour
Stakeholders:	
Description Hi all. Just occurred to me that we could easily remove the need for the operators* to specify the location of the config file. The create_* script(s) have the abspath of the config file when creating the initial work tree... and that config file abspath typically does not change for the lifetime of the work tree (at least n 90%+ cases). So why not use this information as the default argument for the reconfig option? Would make life much easier for the operators. Thanks, Igor • both factory and frontend	
Subtasks: Feature # 2413: Changing the location of setup scripts factory.sh and frontend.sh Assigned	

History

#1 - 11/18/2011 12:06 PM - Igor Sfiligoi

Maybe I have an easier solution;
rename \$WORKDIR/glideinWMS.xml into \$WORKDIR/glideinWMS.current.xml
and symlink \$WORKDIR/glideinWMS.xml -> \$CONFIGDIR/glideinWMS.xml

This way we get the best of both worlds;
a dedicated config dir (which in my opinion is needed, because the XML file is not the only thing a frontend admin needs),
and the "standard" look and feel of reconfig.

How does this sound?

Igor

PS: Same thing for the frontend, of course.

PS2: I would keep the option of providing the actual config file for backwards compatibility.

#2 - 12/12/2011 03:48 PM - Burt Holzman

- Assignee set to Krista Larson

#3 - 12/14/2011 09:51 AM - Parag Mhashilkar

- Target version changed from v2_5_4 to v2_5_5

#4 - 12/14/2011 08:58 PM - Parag Mhashilkar

From Igor ... 2 different emails ...

I am surprised the RPM can work without this.
But I have not looked into the RPM yet, so will comment more in detail later today.

Just so you know what's on my mind,
Igor

Yes, the RPM pretty much is doing this.

Can you check with Doug how much of the RPM stuff is related to this?

#5 - 12/14/2011 10:19 PM - Douglas Strain

Hi, the RPM works by substituting its own "frontend_startup" that passes the (hard-coded) config file to reconfig_frontend. There is no code change, just a script change.

I'm not sure I like the symlink solution proposed, as it seems confusing. If I remember right, the create scripts create the frontend_startup script, so couldn't they "hard-wire" in a default to use if it is not explicitly provided? This is just personal preference though, I am sure I can change the RPM to work with however this turns out.

#6 - 01/19/2012 04:11 PM - Krista Larson

I'm not sure I like the symlink either. We already have multiple versions of the cfg file in these dirs and adding a symlink with the same name seems to add to the confusion (at least for me).

The factory_startup script doesn't know the instance name and I can't use the config file to get the instance name to build the cfg path. So Doug's suggestion of hardcoding the default location in the startup script seems reasonable.

What does everyone else think?

#7 - 01/19/2012 04:31 PM - Parag Mhashilkar

In factory_startup you have access to

```
factory_dir=/local/home/gfactory/master/glideinsubmit/glidein_v1_0
glideinWMS_dir=/opt/glideinwms
```

```
Following gives you instance name
grep glidein $factory_dir/glideinWMS.xml | grep glidein_name | awk -F'glidein_name=' '{print $NF}' | awk -F' ' '{print $1}'
```

And with the two symlinks options above you got access to everything. What else do you need?

#8 - 01/19/2012 04:59 PM - Krista Larson

I guess I was under the impression we should never use the files in the work_dir directly (assuming no symlink). I thought Doug's suggestion of just using a hard-coded value in the file was good and he said it would work with the RPM (he'd adjust his startup script to match).

The above is direct for the factory, but the instance name is not in the frontend xml file. You'd have to parse it from the frontend name in the config. If we allow "-" in the instance or service name this may be confusing. But all this is available on create so it would be easy to add to the startup scripts.

It didn't seem like we had agreed on adding symlinks and I'm not sure what we gain by adding it? If it's only for this, it seems unnecessary and we'd still have to parse the name in the frontend config.

#9 - 01/19/2012 05:06 PM - Douglas Strain

You have to give the frontend.xml as a parameter when you create the glideins / frontend, so why not just take it and put it in the frontend_startup/factory_startup (since you create that file when you create the glideins anyway). If the location changes, the user can always explicitly pass the location. This is not a complicated fix and requires no crazy parsing or symlinks.

#10 - 01/19/2012 05:24 PM - Parag Mhashilkar

Putting the instance/glidein_name and the actual config file in the startup scripts is fine. But with that you are hiding the details from the user which can potentially be useful.

Now user needs to explicitly specify the cfg file so he knows which file he is using. If you are putting this info in the startup scripts (without symlinks) he has to refer to the startup script to remind him which configuration file to use.

#11 - 01/20/2012 11:12 AM - Krista Larson

This seems to just be a shortcut for admins so I think I am missing something. I don't understand what we are hiding or how the symlink helps. What do we do about the bad symlink if the user doesn't use the standard location? Shouldn't they refer to the docs for the startup script behavior (including writeback, files changed, etc)? We could also output a msg stating what config it is using (as well as the other behavior) to reduce confusion.

There still is the problem with the frontend config not containing the instance name - this is only known in the create step. Changing the startup script for this would also break "upgrade" on a frontend. Maybe we should add instance name to the frontend config?

#12 - 01/23/2012 01:34 AM - Igor Sfiligoi

Hi Krista.

Users referring to docs? Have you ever seen a user using the documentation? ;)

Seriously though, I don't see why you insist on "frontend config not containing the instance name"; of course the config does not contain the instance dir... it has no need to.

The proposal was for the `creage_frontend/creat_glidein` to use the `abspath` as passed to it as an argument to figure out where the default frontend/factory xml file was to be located. And reference it at install time (either via symlink or hardcoding)

As for symlink vs hardcoding in the startup script; they both have advantages and disadvantages.

Symlink:

More "standard"... the xml file local to the startup script is actually the right one.

However, this indeed breaks our convention that nothing in the work area should ever be changed by hand.

Hardcoding:

We keep the current semantics (never change anything in the work area), but we indeed hide everything from the user.

Thinking about it, one way out (in the hardcoding option) would be to add a "-config" option to the startup script to tell the user where the config file is. This way we get the best of both worlds.

What do you think?

BTW: This would also allow us to use the stock startup script in the RPM, by just symlinking the work area one into the `/etc/init.d/`.

BTW2: In the "tarnball mode", we could symlink the work area startup script into the instance directory, so the users never actually see the work area... this may create for way less confusion than what we have now... i.e. we would be much more compatible with the "RPM mode". (while still being backwards compatible for for a smooth transition and for those who liked the old operation mode)

Yes... I am starting to like the hardcoding++

Igor

#13 - 01/23/2012 10:23 AM - Parag Mhashilkar

Igor correct me if I am wrong or missed something. Just to make sure what Igor is proposing -

- 1) Move `factory_startup` from `workdir` to the `.cfg` dir. Create a link in the `workdir` for backward compatibility
- 2) Hard code the location of the `*.cfg/glideinwms.xml` in the startup script. This is only required for backward compatibility. Instance/glidein name doesn't really matter now.
- 3) Optionally provide `-config` option to startup scripts that prints the location of config files in `cfg` area (and `workdir` maybe?)
- 4) Couple these changes with [#2413](#)

So everything the user needs is in `.cfg` directory.

Same thing applies to frontend.

#14 - 01/23/2012 10:42 AM - Igor Sfiligoi

- 1) No;
the authoritative `factory_startup` should stay in the work dir; one should never assume anything outside this area is needed to operate the factory (outside the `reconfig` command).

But we do want to create the symlink to that in the instance area, as a convenience for the admins.

BTW: This should be optional in `create_glidein`... i.e. we should support a `--no-startup-symlink` option, for those admins that do not have the initial xml file in a dedicated dir.

- 2) Almost.

It should read

"Hard code the location of input the xml file (i.e `abspath(argv1)`) in the startup script"

It does not need to be `*.cfg/glideinWMS.xml`, although it will likely be, if using the installer.

- 3)

No need to provide the `workdir`... that IS in the xml file.

The `-config` would just point to the source XML file.

BTW: Should we support `--config-persistent` (or similar) that would rewrite the `factory_startup` (in the `workdir`) to the new `abspath`?

Igor

#15 - 01/23/2012 10:56 AM - Parag Mhashilkar

Igor Sfiligoi wrote:

- 1) No;
- 3)

No need to provide the workdir... that IS in the xml file.

The -config would just point to the source XML file.

BTW: Should we support --config-persistent (or similar) that would rewrite the factory_startup (in the workdir) to the new abspath?

I don't understand the use case for --config-persistent

#16 - 01/23/2012 12:09 PM - Igor Sfiligoi

--config-persistent would change the hardcoded config path to the config file used during this reconfig.

The use case is as follows:

At the beginning of time, the user used config /mydir1/gf.xml to create the factory, so this has been hardcoded in the startup script.

3 months down the road, he decides he wants to move the config file to /mydir2/gw.xml. It would be nice if he could get the startup script to use this new location from now on.

Hope this explains,
Igor

PS: --config would instead be for a one-off changes in the config file location.

#17 - 01/23/2012 04:08 PM - Parag Mhashilkar

This is what I thought but wanted to make sure I did not miss anything. Said that, its usefulness is subjective but I won't fight it. This seems more like stripped down version of clone glidein. Option name --relocate-config seems intuitive than --config-persistent

#18 - 01/31/2012 09:36 AM - Krista Larson

- Target version changed from v2_5_5 to v2_5_6

#19 - 03/08/2012 04:14 PM - Krista Larson

- Status changed from New to Resolved

Added this feature the factory and frontend in both v2plus and in master.

#20 - 03/12/2012 09:39 PM - Parag Mhashilkar

- Target version changed from v2_5_6 to v2_5_7

#21 - 04/13/2012 11:39 AM - Parag Mhashilkar

- Target version changed from v2_5_7 to v2_7_x

#22 - 06/12/2012 10:20 AM - Parag Mhashilkar

- Target version changed from v2_7_x to v2_6

#23 - 07/24/2012 10:53 AM - Parag Mhashilkar

- Status changed from Resolved to Closed