Systemd and buffered output

January 28, 2015

Kenny MacDonald submitted LCFG bug #799 reporting that the output of updaterpms appeared to be buffered at boot time. He is correct, and this behaviour is particularly noticeable when installing a large-ish number of substantial RPMs – it is not at all obvious that the system is doing anything productive. There is a real danger that a desktop user may decide to power-cycle/reboot his/her desktop should they encounter this behaviour.

By default, systemd units are configured to output to '*journal+console*'. This means that output is sent to both journald and to */dev/console*. For some reason, with this configuration, output to */dev/console* is buffered (multiple lines). If a unit is configured to output to '*tty*' (which is equivalent to '*console*'), lines are only singly buffered. I can only guess that systemd is blocking on logging to journald, and the output to */dev/console* is blocked behind this.

The only obvious solution is for lcfg-updaterpms to output just to '*tty*'. This does mean that we lose the ability to log useful error messages that are produced during the updaterpms run (by install scripts), but we haven't captured those error messages in the past.

The current solution (work-around, really) is to output to 'tty' by default, but allow buffered output to be enabled by use of

#define LCFG_UPDATERPMS_BUFFERED_OUTPUT

at the top of a machine's profile. Buffered output could usefully be enabled on machines used for release testing.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

Systemd – LCFG documentation

December 5, 2014

A start has been made on LCFG systemd documentation.

The Systemd Cookbook documents common queries / recipes.

The lcfg-systemd document is a more full description of LCFG systemd.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

Systemd – fun with targets dependencies

December 5, 2014
If you have a target B.target which wants/requires S.service and you want S.service to start after A.target – it isn't sufficient to state B.target 'requires' and 'after' A.target. If you just do that, S.service will start based on its own requires/afters – usually much earlier than after A.target. You also need to state that S.service 'requires' and 'after' A.target.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

Systemd – issues with overriding dependencies

December 5, 2014

It does not appear possible to override dependencies defined through a unit's .wants or .requires directories.

A unit can declare a 'want' or 'require' either through directives in its unit file or through links created in a .wants or .requires directory.

You can easily override dependencies specified in a unit's unit file, by creating a file of the same name in /etc with the offending directives removed. This file is simply used instead of the file in /lib. For example, unit sshd.service might specify that it 'Wants' kerberos.service by means of a Wants directive in its unit file /lib/systemd/sshd.service. One can override this by creating an /etc/systemd/sshd.service with the offending Wants directive removed.

One might expect that creating unit's .wants or .requires directories in /etc would override the equivalent directories in /lib, but that is not the case. For example, ntpd.service may specify that it 'Requires' network.target by means of a symlink, in /lib/systemd/ntpd.service.requires, to the network.target unit file. One might expect that you could remove this dependency by creating a directory /etc/systemd/ntpd.service.requires without the offending symlink. Unfortunately this isn't what happens – instead, systemd uses the union of /etc/systemd/ntpd.service.requires and /lib/systemd/ntpd.service.requires.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

Systemd gains start_on_add resource

September 23, 2014

The systemd component has gained the ability to auto-start service units when they are added to the units resource, without requiring a reboot. Setting the start_on_add_{tag} resource to true will enable this

functionality for the unit associated with {tag}.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

Systemd presets

September 18, 2014

Once we had managed to install EL7 machines using the LCFG installroot, we discovered that the systemd config for nslcd, which had been created by the LCFG systemd component whilst running the installroot, was removed at the first reboot. On investigation we discovered that when the nslcd RPM was being installed, the RPM postinstall script was calling 'systemctl preset'.

Systemd preset files (description) are used to encode policy as to which units should be enabled by default and which should be disabled. The preset files, which live under /usr/lib/systemd/system-preset, are provided by the distribution.

Redhat's default preset policy for nslcd is 'disabled', so when 'systemctl preset' was called as part of the RPM install, systemd removed the nslcd configuration that the LCFG systemd component had previously created.

Fortunately you can disable individual preset files by creating a softlink (to /dev/null) with the same filename in /etc/systemd/system-preset.

In the LCFG world, we'll have no use for the preset mechanism as we'll be using the LCFG systemd component to control which systemd units are enabled. The LCFG systemd component can be configured to disable individual preset files.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

Fun with systemd targets

September 18, 2014

The existing LCFG boot component allows one to configure components such that they are not started, at system boot time, until after those components which may ask for a system reboot have completed. Simply put, the boot component starts all components with a start ID < 100. It then reboots the system if any of those components have requested a reboot. If no reboot has been requested, it then starts those components with a start ID >= 100. This is so that user fronting services such as web, getty and ssh aren't started just for them to be taken down should a reboot be requested.

We wanted to replicate this behaviour with systemd. It appeared obvious(?) that targets would assist us in doing this.

Our first attempt was to create a new target lcfg-multi-user-stable.target to be the end 'default' target, and have this require the normal multi-user.target. The following diagram shows the basic structure :-

- lcfg-multi-user.target (as end 'default' target)
  - multi-user.target

        - lcfg-auth.service
        - lcfg-client.service
        - ….
        - lcfg-updaterpms.service
        - getty.service
        - other stock units
    - lcfg-rebootcheck.service (after multi-user.target)
    - lcfg-clearbootctx.service (after multi-user.target)
    - lcfg-cron.service (after multi-user.target)
    - lcfg-apacheconf.service (after multi-user.target)
    - lcfg-openssh.service (after multi-user.target)

The lcfg-rebootcheck tool reboots the system if any component started whilst reaching multi-user.target has requested a reboot.

The above configuration appeared to work fine, but for the fact that getty was being started before the reboot check was being performed. From the above table, we can see that we want to start getty.service after multi-user.target (or lcfg-rebootcheck.service). Unfortunately it proved impossible to achieve this. Getty is, by default, pulled in as a "wants" of multi-user.target. This dependency is hard-coded in /usr/lib/systemd. Amazingly, whilst it is possible to override a unit file in /usr/lib/systemd/system by creating a file with the same name in /etc/systemd/system, or to disable a unit file by creating a symlink of the same name to /dev/null, it is not possible to override dependency files in .wants or .requires directories.

For our second attempt, we decided to keep the default multi-user.target as the end 'default' target, adding new LCFG targets as "wants" of multi-user.target.

- multi-user (as end 'default' target)
    - lcfg-multi-user-stable
        - lcfg-multi-user
            - lcfg-auth
            - lcfg-client
            - …
            - lcfg-updaterpms
        - lcfg-rebootcheck (after lcfg-multi-user)
    - lcfg-clearbootctx (after lcfg-multi-user-stable)
    - lcfg-cron (after lcfg-multi-user-stable)
    - lcfg-apacheconf (after lcfg-multi-user-stable)
    - lcfg-openssh (after lcfg-multi-user-stable)
    - getty (after lcfg-multi-user-stable)

The above configuration achieves what we wanted.

Whilst we have worked round the inability to override dependencies in this case, we're convinced we'll trip up over this problem again in the future.

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

LCFG and systemd

February 11, 2014

Alastair gave a quick introduction to *systemd* to the Informatics COs and explained some of the problems we're going to face with integrating LCFG with *systemd*.

In summary, to integrate LCFG with *systemd* properly will require some fundamental changes to the LCFG framework. Unfortunately, we do not have time to properly consider such changes before we have to deploy RHEL 7 desktops, so we have no option but to find some way of using the existing framework with *systemd*. It's likely that this will involve taking different approaches for different components. Once we have finished the desktop port, we will take stock and consider what changes need to be made to the LCFG framework – before we start work on the server port.

(Alastair's slides are here)

Leave a Comment » | systemd | Permalink
Posted by Alastair Scobie

- You are currently browsing the archives for the systemd category.
- Search Search

- **Recent Posts**
  - internet-online.target
  - New installer feature
  - openafs client changes
  - Service function improvements
  - Dynamic image generation for login screens

- **Recent Comments**
  - Service function improvements | Scientific Linux 7 LCFG port diary on ngeneric changes
  - grub2 – sanity prevails | Scientific Linux 7 LCFG port diary on grub2
  - squinney on Bugs
  - Kenny MacDonald on Bugs

- **Archives**
  - March 2015
  - February 2015
  - January 2015
  - December 2014
  - November 2014
  - October 2014
  - September 2014
  - August 2014
  - July 2014
  - June 2014
  - April 2014
  - March 2014
  - February 2014
  - January 2014
  - December 2013

- # Categories

  - systemd
  - Uncategorized

- # Meta

  - Site Admin
  - Log out
  - Entries RSS
  - Comments RSS
  - WordPress.org

Theme: Contempt by Vault9.