datto

Golden Images for Scaling Up with the Best of Them

Neal Gompa

Who am I?

- Professional technologist
- Linux user for nearly fifteen years
- Contributor and developer in Fedora, CentOS, openSUSE, Mageia, and OpenMandriva Linux distributions
- Contributor to RPM, DNF, KIWI, and various related projects
- Senior DevOps Engineer at Datto, Inc.

The beginning of the Datto Cloud

Pre-automation era

In the beginning of the Datto Cloud, there was not much in the way of automation for provisioning systems. Each physical server was built and the operating system was installed by hand by going through the installation process manually.

This worked when we had few physical servers (and few employees doing the work), but quickly became a problem as we grew.

Minimal automation era

Eventually, we grew to a point where each system being subtly different due to the manual installation of the operating system and server software became a problem.

We introduced installation automation (kickstart/kickseed/pre-seed) to regularize the system software installation process.

Virtual machines on physical hosts were provisioned similarly.

The Foreman era

Bringing in the Foreman

The need to further standardize and automate system deployments necessitated introducing configuration management and lifecycle management system.

Thus, the **Foreman** with **Puppet** was introduced to automate the configuration and maintain the standard configuration as it changed centrally.

Puppet-master era

As we introduced Foreman, we started using it as part of provisioning virtual machines.

At first, we kept the same installation process and then auto-connected the VM to Foreman to run Puppet. Once we started running an OpenStack system, we started using official golden images and then running Puppet on there to speed things up considerably.

Puppet-master era

Eventually, our Puppet became so complex that Puppet runs were taking upwards to an hour for initial runs on images.

We started splitting up our Puppet manifests and leveraging Packer to pre-bake “common” configuration, while still running application-specific stuff at provision-time.

Unfortunately, this did not scale well as more products and teams needed to work with it using other tools in a self-service manner.

Into the era of self-service

What started going wrong…

The workflow we were using to build our images worked great… up to a point. The images were largely controlled by the infrastructure team and the content and nature of the images made it difficult for software engineers to influence them for their needs.

Furthermore, a gradual shift away from Puppet started, in line with a shirt toward software engineering teams owning more of the operational nature of their products and services.

The new requirements…

Integrating new products and teams meant we needed to rethink how our cloud images were made for them to use. This was distilled into the following new requirements:

- Multi-distro (CentOS and Ubuntu)
- Agnostic and independent of configuration management tools
- Unified system tooling and interfaces across distributions (as much as possible)
- Corporate standard tools integrated into the baseline for consumers

Rethinking the image build

… with some kiwi?

Searching for a new image build tool

As it turns out, when you need unified tooling that supports multiple distribution families, the list of viable options are quite short.

Even without that, a lot of build tools are purpose-built, or made and then get no maintenance. Worse yet, most of these tools have little to no community development around them.

Selecting KIWI

After a fair bit of searching, it came down to two options:

- mkosi
- KIWI

We selected KIWI primarily because of its maturity and stronger community. In particular, the input manifest format and SBOM logs it creates as part of the image build made it much more attractive.

Selecting KIWI

- Straightforward and idiomatic
  - XML/YAML/JSON descriptions with script hooks
- Flexible
  - Builds almost any type of image
  - Provides an API to construct custom image types
- Automatically produced SBOM artifact logs
- Free and Open Source Software (GPLv3+)
- Actively developed and maintained
- Friendly developers

Beginning our use of KIWI

Improving KIWI

Once we settled on KIWI, we started trying to adapt some of our image builds to use it and came across a few issues we needed to resolve to make it fully ready for our use.

So, we rolled up our sleeves and contributed improvements!

Improving KIWI

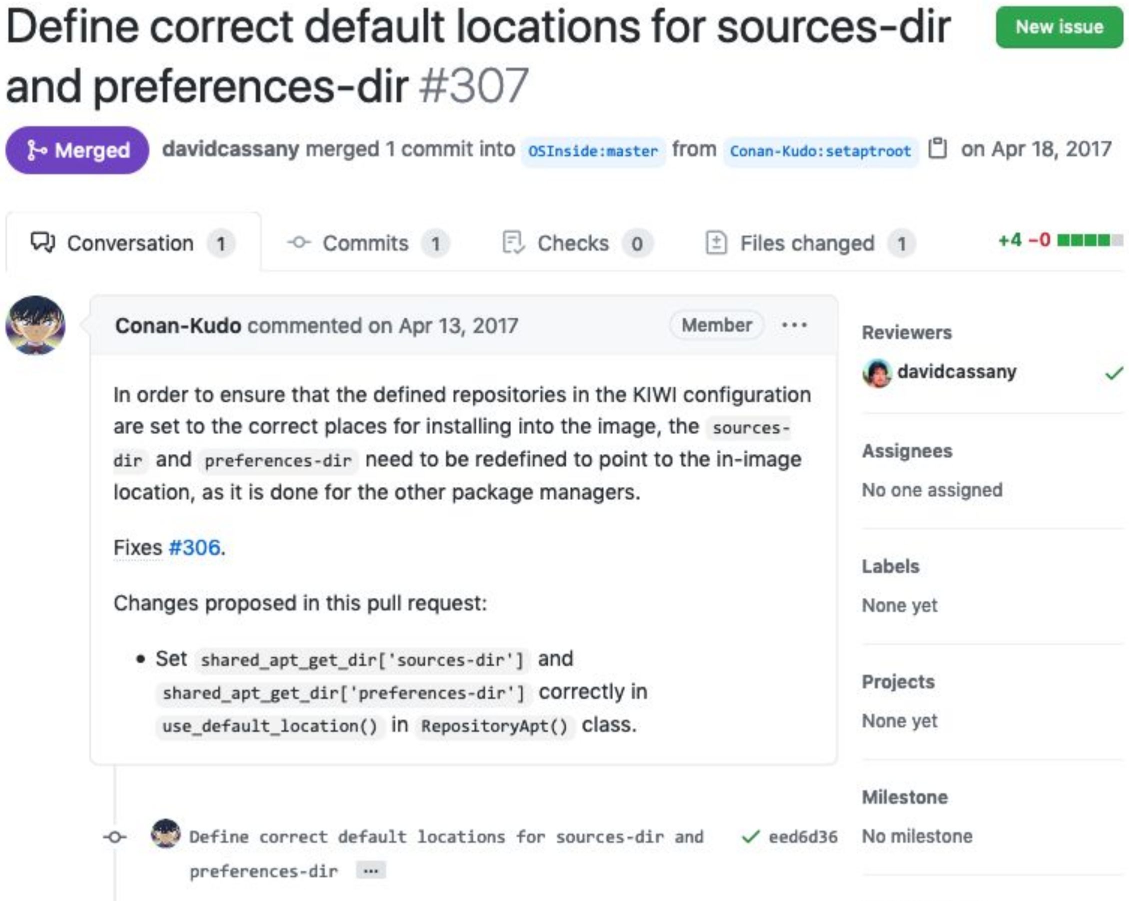

Define correct default locations for sources-dir and preferences-dir #307

Merged davidcassany merged 1 commit into `OSInside:master` from `Conan-Kudo:setaptroot` on Apr 18, 2017

Conan-Kudo commented on Apr 13, 2017 Member

In order to ensure that the defined repositories in the KIWI configuration are set to the correct places for installing into the image, the `sources-dir` and `preferences-dir` need to be redefined to point to the in-image location, as it is done for the other package managers.

Fixes #306.

Changes proposed in this pull request:

- Set `shared_apt_get_dir['sources-dir']` and `shared_apt_get_dir['preferences-dir']` correctly in `use_default_location()` in `RepositoryApt()` class.

```
Define correct default locations for sources-dir and ✓ eed6d36
preferences-dir ...
```

Reviewers
davidcassany ✓

Assignees
No one assigned

Labels
None yet

Projects
None yet

Milestone
No milestone

Improving KIWI

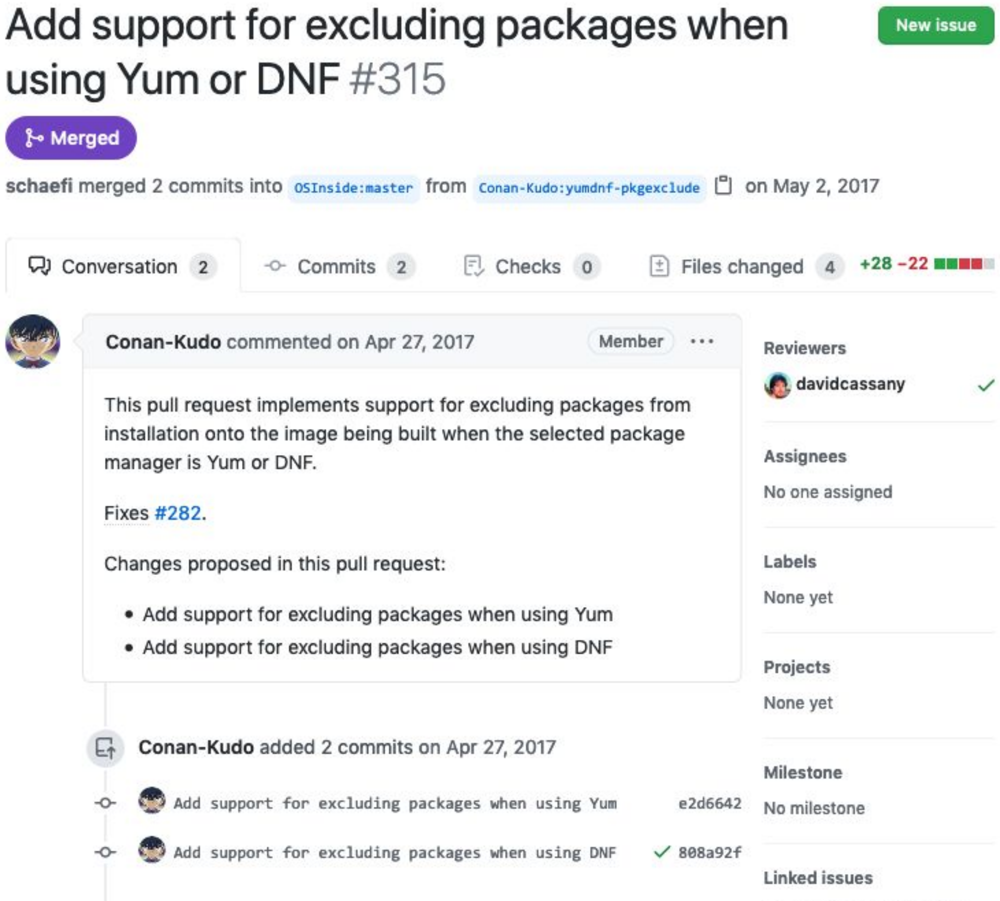

Add support for excluding packages when using Yum or DNF #315

schaefi merged 2 commits into `OSInside:master` from `Conan-Kudo:yumdnf-pkgexclude` on May 2, 2017

Conan-Kudo commented on Apr 27, 2017 Member

This pull request implements support for excluding packages from installation onto the image being built when the selected package manager is Yum or DNF.

Fixes #282.

Changes proposed in this pull request:

- Add support for excluding packages when using Yum
- Add support for excluding packages when using DNF

Conan-Kudo added 2 commits on Apr 27, 2017

```
Add support for excluding packages when using Yum    e2d6642
Add support for excluding packages when using DNF  ✓ 808a92f
```

Reviewers: davidcassany ✓

Assignees: No one assigned

Labels: None yet

Projects: None yet

Milestone: No milestone

Linked issues

datto

Improving KIWI

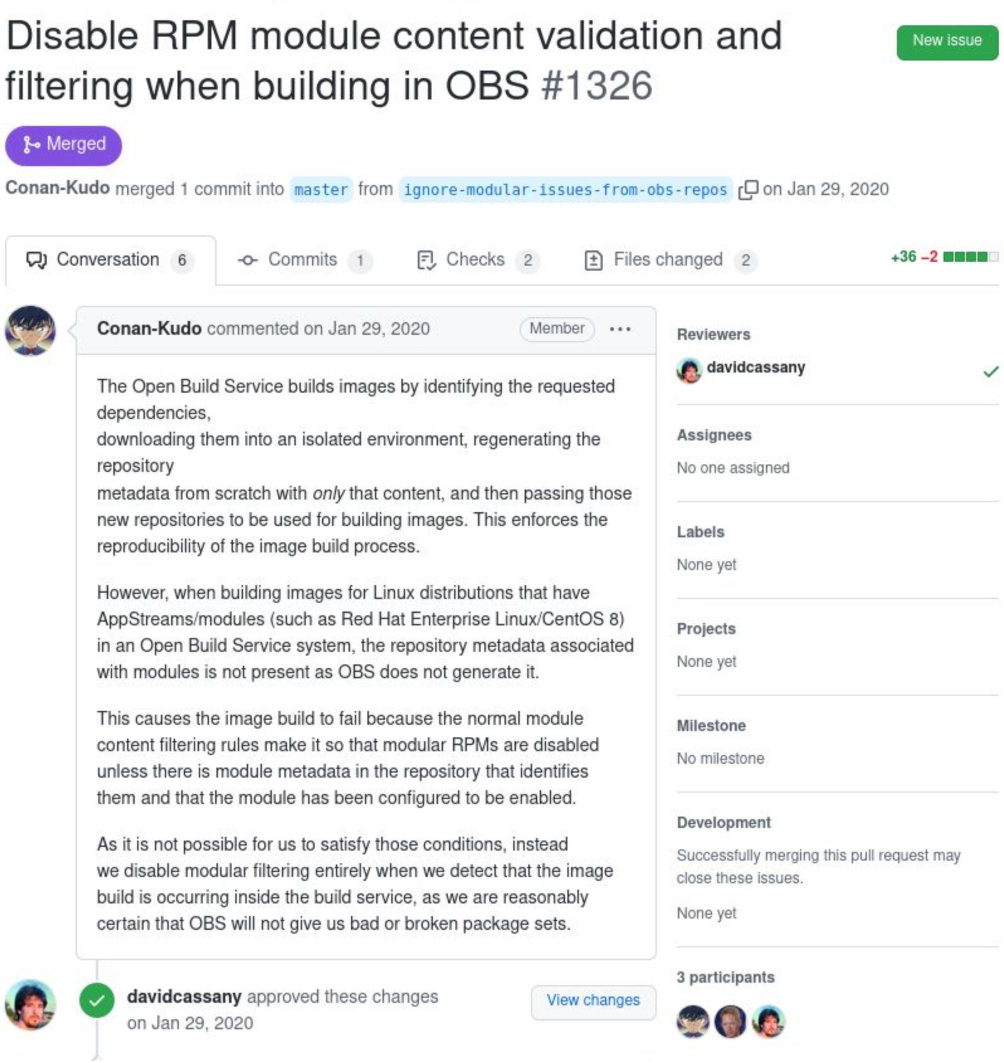

Disable RPM module content validation and filtering when building in OBS #1326

Conan-Kudo merged 1 commit into `master` from `ignore-modular-issues-from-obs-repos` on Jan 29, 2020

Conan-Kudo commented on Jan 29, 2020 Member

The Open Build Service builds images by identifying the requested dependencies,
downloading them into an isolated environment, regenerating the repository
metadata from scratch with *only* that content, and then passing those new repositories to be used for building images. This enforces the reproducibility of the image build process.

However, when building images for Linux distributions that have AppStreams/modules (such as Red Hat Enterprise Linux/CentOS 8) in an Open Build Service system, the repository metadata associated with modules is not present as OBS does not generate it.

This causes the image build to fail because the normal module content filtering rules make it so that modular RPMs are disabled unless there is module metadata in the repository that identifies them and that the module has been configured to be enabled.

As it is not possible for us to satisfy those conditions, instead we disable modular filtering entirely when we detect that the image build is occurring inside the build service, as we are reasonably certain that OBS will not give us bad or broken package sets.

davidcassany approved these changes on Jan 29, 2020

Reviewers
davidcassany ✓

Assignees
No one assigned

Labels
None yet

Projects
None yet

Milestone
No milestone

Development
Successfully merging this pull request may close these issues.
None yet

3 participants

Improving KIWI

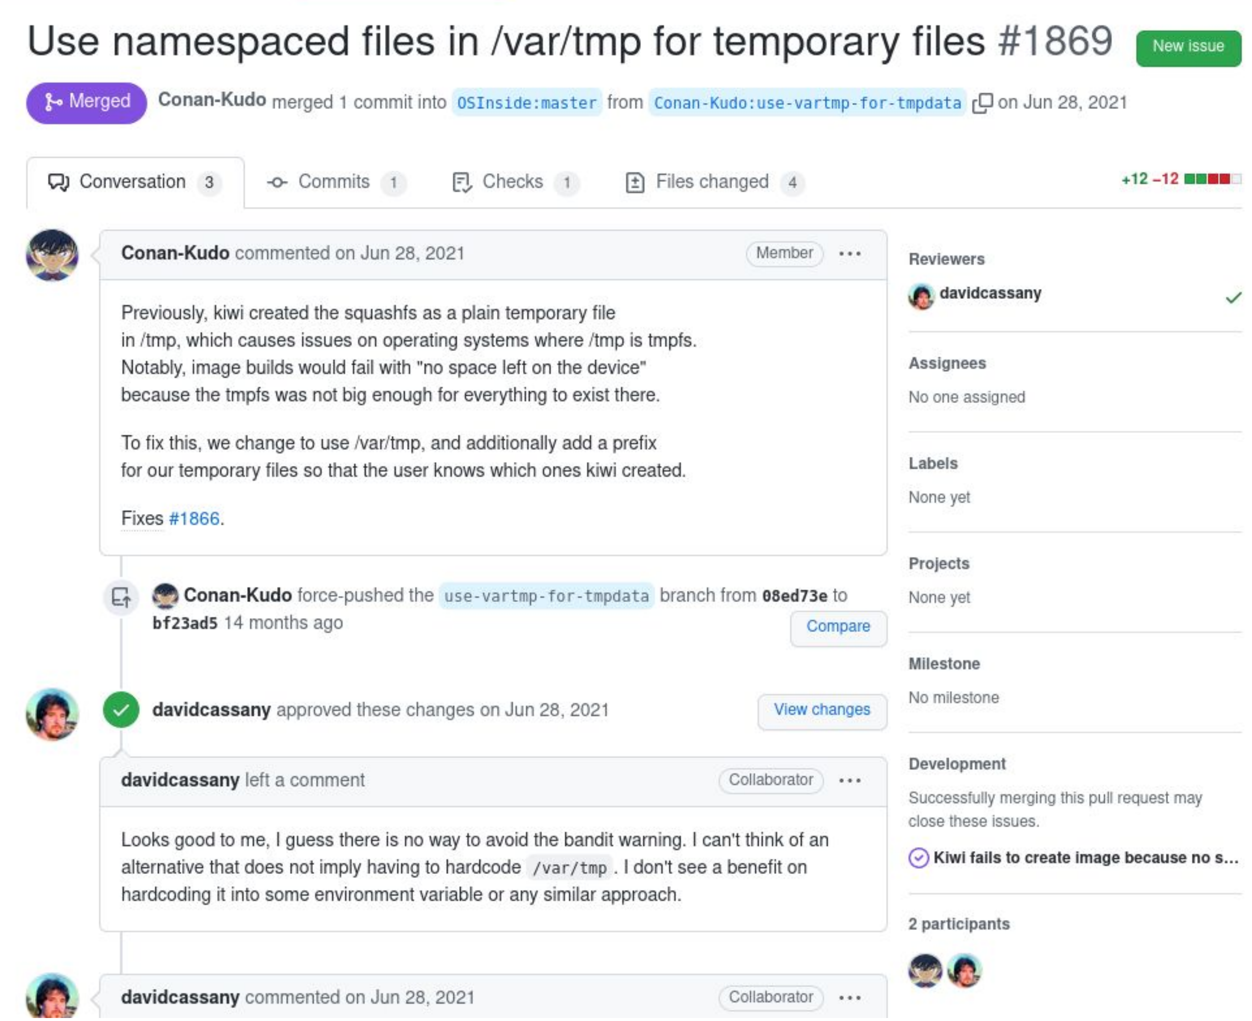

Use namespaced files in /var/tmp for temporary files #1869

Conan-Kudo merged 1 commit into `OSInside:master` from `Conan-Kudo:use-vartmp-for-tmpdata` on Jun 28, 2021

Conan-Kudo commented on Jun 28, 2021 Member

Previously, kiwi created the squashfs as a plain temporary file in /tmp, which causes issues on operating systems where /tmp is tmpfs. Notably, image builds would fail with "no space left on the device" because the tmpfs was not big enough for everything to exist there.

To fix this, we change to use /var/tmp, and additionally add a prefix for our temporary files so that the user knows which ones kiwi created.

Fixes #1866.

Conan-Kudo force-pushed the `use-vartmp-for-tmpdata` branch from **08ed73e** to **bf23ad5** 14 months ago

davidcassany approved these changes on Jun 28, 2021

davidcassany left a comment Collaborator

Looks good to me, I guess there is no way to avoid the bandit warning. I can't think of an alternative that does not imply having to hardcode `/var/tmp`. I don't see a benefit on hardcoding it into some environment variable or any similar approach.

davidcassany commented on Jun 28, 2021 Collaborator

Reviewers
davidcassany ✓

Assignees
No one assigned

Labels
None yet

Projects
None yet

Milestone
No milestone

Development
Successfully merging this pull request may close these issues.

Kiwi fails to create image because no s...

2 participants

Improving KIWI

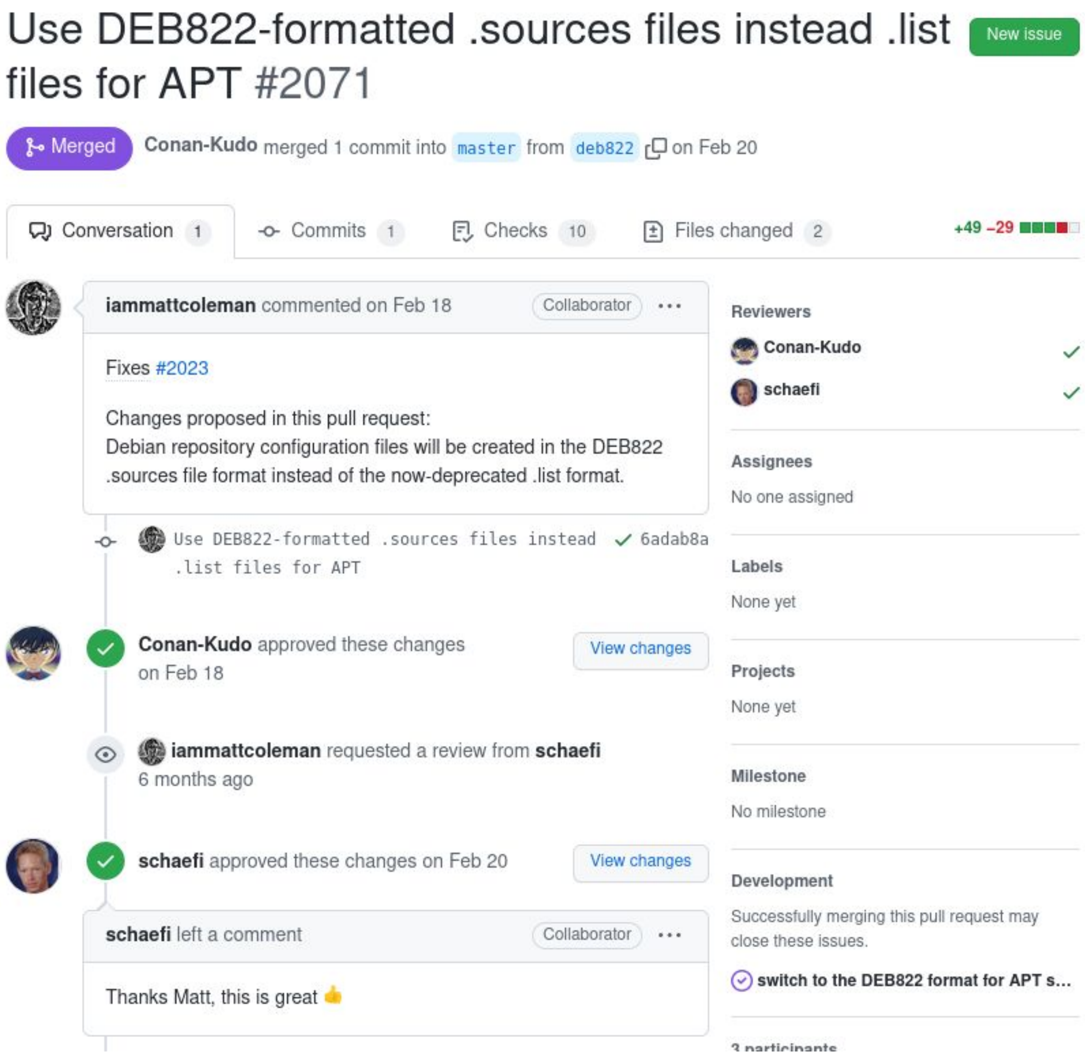

Use DEB822-formatted .sources files instead .list files for APT #2071

Merged **Conan-Kudo** merged 1 commit into `master` from `deb822` on Feb 20

iammattcoleman commented on Feb 18 · Collaborator

Fixes #2023

Changes proposed in this pull request:
Debian repository configuration files will be created in the DEB822 .sources file format instead of the now-deprecated .list format.

```
Use DEB822-formatted .sources files instead ✓ 6adab8a
.list files for APT
```

Conan-Kudo approved these changes on Feb 18 · View changes

iammattcoleman requested a review from **schaefi** 6 months ago

schaefi approved these changes on Feb 20 · View changes

schaefi left a comment · Collaborator

Thanks Matt, this is great 👍

Reviewers
Conan-Kudo ✓
schaefi ✓

Assignees
No one assigned

Labels
None yet

Projects
None yet

Milestone
No milestone

Development
Successfully merging this pull request may close these issues.
switch to the DEB822 format for APT s...

3 participants

Demonstration

datto

Our production pipeline

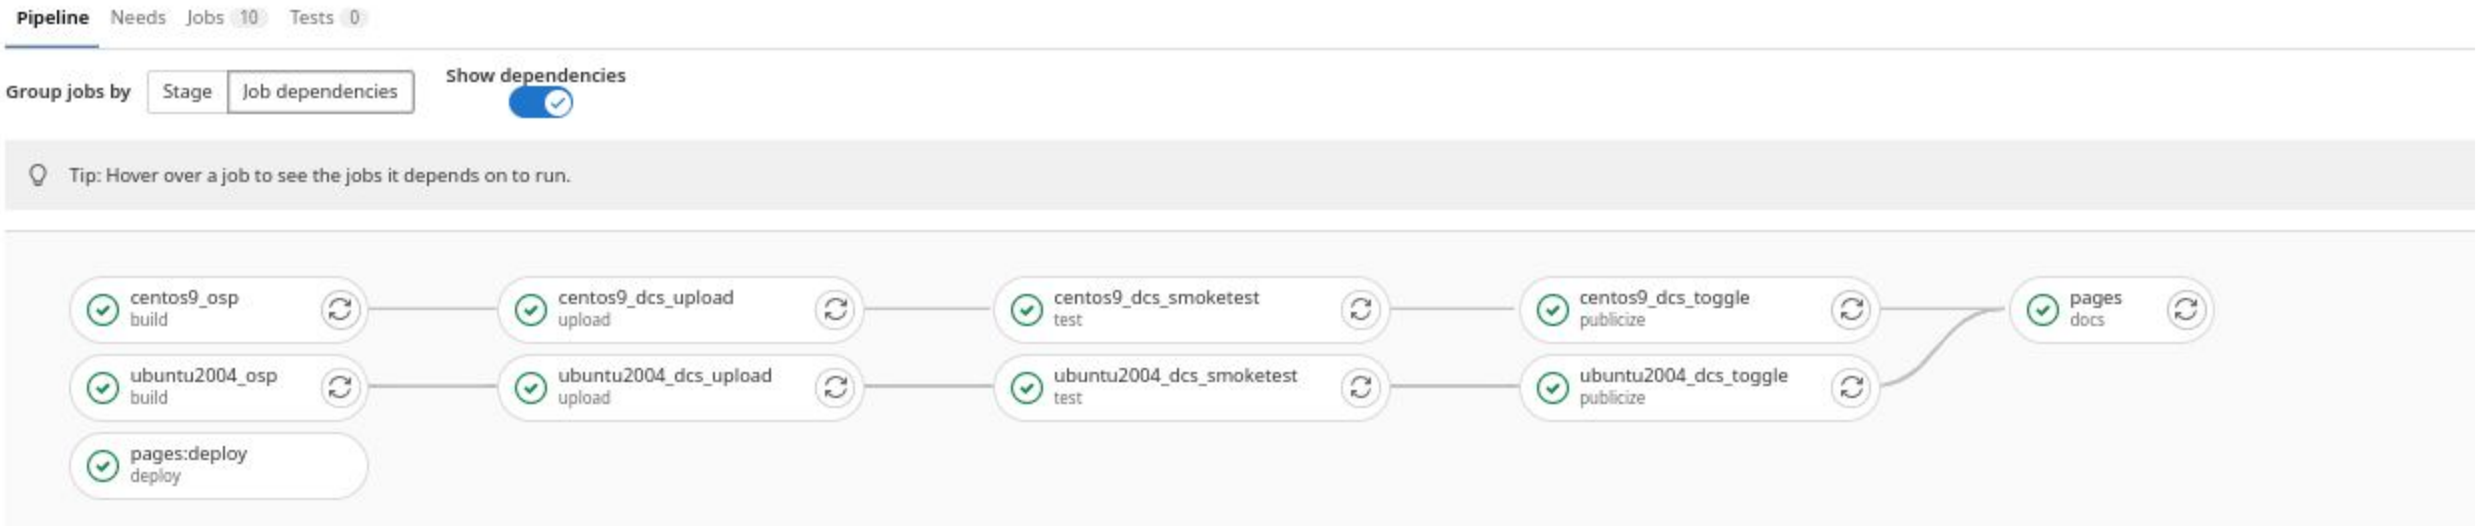

Tip: Hover over a job to see the jobs it depends on to run.

References

- KIWI website and docs: http://osinside.github.io/kiwi/
- KIWI GitHub project: https://github.com/OSInside/kiwi
- Sample descriptions: https://github.com/OSInside/kiwi-descriptions
- Demo descriptions: https://github.com/datto/devconfus22-demo-golden-image-descriptions

datto

The world's leading provider of MSP-delivered IT solutions

Blog - datto.engineering

Careers - datto.com/careers

GitHub - github.com/datto

GitLab - gitlab.com/datto