



Cloud-friendly APT repository publishing

Contents

Why we need a new APT publisher	2
Alternatives to <code>reprepro</code>	3
Aptly	3
Pulp	4
Conclusion	5
Implementation plan	6

Why we need a new APT publisher

Apertis relies on [OBS](https://apertis-website-0b3586.pages.apertis.org/architecture/distribution/workflow-guide/)¹ for building and publishing binary packages. However, upstream OBS provides an APT publisher based on `dpkg-scanpackages`, which is not suitable for a project the scale of Apertis, where a single OBS project contains a lot of packages.

Therefore, our OBS instance uses a custom publisher based on `reprepro`, but it is still subject to some limitations that are now more noticeable as the scale of Apertis has grown considerably:

- It purely acts on the events created by OBS, which means that if they do not get successfully processed immediately the repository will go out-of-sync
- The repositories are hosted on the same machine as OBS, and there is no way for external tools to interact with them; for instance there is no way to request the creation of snapshots with a different strategy than the current one of creating a snapshot for every single change
- When branching a release `reprepro` has to be invoked manually to initialize the exported repositories
- When branching a release the OBS publisher has to be manually disabled or it will cause severe lock contention with the manual invocation mentioned above
- Removing a package requires manual intervention
- Snapshots are not supported natively
- Cloud storage is not supported

In order to address these shortcomings, we need to develop a new APT publisher (based on a backend other than `reprepro`) which should be capable of:

- Publishing the whole Apertis release on non-cloud storage
- Publishing the whole Apertis release on cloud storage
- Natively supporting snapshots

¹<https://apertis-website-0b3586.pages.apertis.org/architecture/distribution/workflow-guide/>

- Automatic branching of an Apertis release, not requiring manual intervention on the APT publisher
- Using a synchronization strategy to ensure that OBS and APT repositories automatically tend to consistent state:
 - removing a package from OBS should trigger the removal of the package from the APT repositories as well
 - once a publishing failure is resolved (network issues, etc.) the publisher should recover automatically

Alternatives to reprepro

The Debian wiki includes [a page](#)² listing most of the software currently available for managing APT repositories. However, a significant portion of those tools cover only one of the following use-cases:

- managing a small repository, containing only a few packages
- replicating a (sometimes simplified) official Debian infrastructure

A few of the mentioned tools, however, are aimed at managing large-scale repositories within a custom infrastructure, and offer more advanced features which could be of interest to Apertis. Those are:

- [aptly](#)
- [pulp](#)

[Laniakea](#)³ was also considered, but as it's meant to work within a full Debian-like infrastructure and doesn't offer any cloud-based storage option, it was dismissed as well.

Extended search did not point to other alternative solutions covering our use-case.

Aptly

[Aptly](#)⁴ is a complete solution for Debian repository management, including mirroring, snapshots and publication.

It uses an internal, locally-stored package pool and database, and provides cloud storage options for publishing ready-to-serve repositories. Aptly also provides a full-featured CLI client and an almost complete REST API. It could therefore run either directly on the same server as OBS, or on a different one. The REST API misses mirroring support for now, so these features can only be used from the command-line client.

Package import and repository publication are separate operations:

²<https://wiki.debian.org/DebianRepository/Setup>

³<https://github.com/lkhq/laniakea>

⁴<https://www.aptly.info/>

- The package is first imported to the internal package pool and associated to the requested repository in a single operation
- When all required packages are imported, the repository can be published atomically

Repositories can be published both to the local filesystem and to a cloud-based storage service (Amazon S3 or OpenStack Swift).

Moreover, Aptly identifies each package using the (name, version, architecture) triplet: by doing so, it allows keeping multiple versions of the same package in a single repository, while `reprepro` kept only the latest package version. This requires additional processing for Aptly to replicate the current behavior.

Finally, attention should be paid to regularly cleaning up the database and package pool: unused packages are kept in the pool, even when obsoleted by a newer version and/or removed from all repositories, until a database cleanup is triggered. A daily cleanup job should be sufficient to make sure the internal pool doesn't carry unused packages over time.

Pros

- tailored for APT repository management: includes some interesting features such as multi-component publishing
- command-line or REST API interface (requires an additional HTTP server for authentication and permissions management)

Cons

- uses a local package pool which can grow large if a lot of packages and versions are used simultaneously
- requires additional processing to keep only the latest version of each package
- needs regular database cleanups

Pulp

`Pulp`⁵ is a generic solution for storing and publishing binary artifacts. It uses plugins for managing specific artifact types, and offers a plugin for DEB packages.

It offers flexible storage options, including S3 and Azure, which can also be extended as the storage backend is built on top of `django-storages`, which provides a number of additional options.

Pulp can be used through a REST API, and provides a command-line client for wrapping a significant portion of the API calls. Unfortunately, the DEB

⁵<https://pulpproject.org/>

105 plugin isn't handled by this client, meaning only the REST API is available for
106 managing those packages.

107 Its package publication workflow involves several Pulp objects:

- 108 • the binary artifact (package) itself
- 109 • a Repository
- 110 • a Publication
- 111 • a Distribution

112 Each Distribution is tied to a single Publication, which is itself tied to a specific
113 Repository version. As each Repository modification increments the Repository
114 version, adding or removing a package involves the following steps:

- 115 • add or remove the package from the Repository
- 116 • retrieve the latest Repository version
- 117 • create a new Publication for this repository version
- 118 • update the Distribution to point to the new Publication
- 119 • remove the previous Publication

120 This workflow feels too heavy and error-prone when working with a distribution
121 the scale of Apertis, where lots of packages are often added or updated. Addi-
122 tionally, each Distribution must have its own base URL, preventing publishing
123 multiple Apertis versions and components in the same repository.

124 Pros

- 125 • generic artifacts management solution: can be re-used for storing non-
126 package artifacts too
- 127 • flexible storage options

128 Cons

- 129 • complex workflow for publishing/removing packages
- 130 • unable to store multiple repositories on the same base URL
- 131 • can only be used through REST API

132 Conclusion

133 Based on the above software evaluation, `aptly` seems to be the more appropriate
134 choice:

- 135 • supports snapshots
- 136 • can make use of both local and cloud-based storage for publishing reposi-
137 tories
- 138 • provides useful features aimed specifically at APT repository management
- 139 • allow publishing several repositories and components to a single endpoint

140 Its main shortcoming (locally-stored package pool) can be addressed by imple-
141 menting an option for storing the pool on cloud-based storage. This would be
142 the most efficient approach when compared to the alternative (hosting `aptly` on
143 a remote server and using it through the REST API).

144 Moreover, the following points must be kept in mind when implementing the
145 publisher:

- 146 • `aptly` doesn't remove previous versions of an updated package; although
147 this behavior could be implemented in `aptly` itself, it will be less effort to
148 have the publisher handle removing obsoleted packages
- 149 • the package pool will keep growing as new and updated packages are
150 added, it should therefore be cleaned up on a regular basis by triggering
151 database cleanups
- 152 • publishing large repositories with `aptly` can take a long time; decoupling
153 the action of adding a package from the actual repository publication
154 would be a useful optimization, however it would be outside the scope of
155 the initial implementation

156 Finally, `aptly` is actively maintained upstream, with a new team of developers
157 having taken over its development last year. The chances of it being abandoned
158 and/or replaced with a different project are therefore very low.

159 **Implementation plan**

- 160 • Update OBS to a more recent upstream version: this will provide a more
161 up-to-date base on which we can develop and upstream the new APT
162 publisher
- 163 • Start with a prototype, local-only version capable of:
 - 164 – adding a package to a (manually created) local repository
 - 165 – publishing the repository to local storage
 - 166 – deleting a package from the repository when removing it from OBS
- 167 • Implement automated branching and repository creation for new OBS
168 projects
- 169 • Automate periodic database cleanups
- 170 • Add configuration options for publishing to cloud-based storage
- 171 • Implement cloud-based storage options for `aptly`'s internal package pool