



Cloud-friendly APT repository publishing

# 1 Contents

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

## 8 Why we need a new APT publisher

9 Apertis relies on [OBS](#)<sup>1</sup> for building and publishing binary packages. However,  
10 upstream OBS uses `dpkg-scanpackages` to publish APT repositories in a simplistic  
11 way, which is not suitable for a project the scale of Apertis, where a single OBS  
12 project contains a lot of packages.

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

- 16 • When branching a release `reprepro` has to be invoked manually to initialize  
17 the exported repositories
- 18 • When branching a release the OBS publisher has to be manually disabled  
19 or it will cause severe lock contention with the manual command above
- 20 • Removing a package requires manual intervention
- 21 • Snapshots are not supported natively
- 22 • Cloud storage is not supported

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

- 25 • Publishing the whole Apertis release on non-cloud storage
- 26 • Publishing the whole Apertis release on cloud storage
- 27 • Automatic branching of an Apertis release, not requiring manual interven-  
28 tion on the APT publisher
- 29 • Synchronize OBS and APT repositories; as an example, removing a pack-  
30 age from OBS should trigger the removal of the package from the APT  
31 repositories as well

## 32 Alternatives to reprepro

33 The Debian wiki includes [a page](#)<sup>2</sup> listing most of the software currently available  
34 for managing APT repositories. However, a significant portion of those tools

<sup>1</sup><https://martyn.pages.apertis.org/apertis-website/architecture/workflow-guide/>

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

35 cover only one of the following use-cases:

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

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

- 41 • [aptly](#)
- 42 • [pulp](#)

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

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

## 48 **Aptly**

49 [Aptly](#)<sup>4</sup> is a complete solution for Debian repository management, including  
50 mirroring, snapshots and publication.

51 It uses a local pool and database and provides cloud storage options for publish-  
52 ing ready-to-serve repositories. Aptly also provides a full-featured CLI client  
53 and an almost complete REST API, only missing mirroring support. It could  
54 therefore run either directly on the same server as OBS, or on a different one.

55 Package import and repository publication are separate operations:

- 56 • The package is first imported to the local pool and associated to the  
57 requested repository in a single operation
- 58 • When all required packages are imported, the repository can be published  
59 atomically

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

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

## 66 **Pros**

- 67 • tailored for APT repository management: includes some interesting fea-  
68 tures such as dependency resolving and multi-component publishing

---

<sup>3</sup><https://github.com/lkhq/laniakea>

<sup>4</sup><https://www.aptly.info/>

- 69 • command-line or REST API interface (requires an additional HTTP server  
70 for authentication and permissions management)

## 71 **Cons**

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

## 77 **Pulp**

78 [Pulp](#)<sup>5</sup> is a generic solution for storing and publishing binary artifacts. It uses  
79 plugins for managing specific artifact types, and offers a plugin for DEB pack-  
80 ages.

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

84 Pulp can be used through a REST API, and provides a command-line client  
85 for wrapping a significant portion of the API calls. Unfortunately, the DEB  
86 plugin isn't handled by this client, meaning only the REST API is available for  
87 managing those packages.

88 Its package publication workflow involves several Pulp objects:

- 89 • the binary artifact (package) itself
- 90 • a Repository
- 91 • a Publication
- 92 • a Distribution

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

- 96 • add or remove the package from the Repository
- 97 • retrieve the latest Repository version
- 98 • create a new Publication for this repository version
- 99 • update the Distribution to point to the new Publication
- 100 • remove the previous Publication

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

---

<sup>5</sup><https://pulpproject.org/>

105 **Pros**

- 106 • generic artifacts management solution: can be re-used for storing non-  
107 package artifacts too
- 108 • flexible storage options

109 **Cons**

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

113 **Conclusion**

114 Based on the previous software evaluation, `aptly` seems to be the more appro-  
115 priate choice:

- 116 • supports snapshots
- 117 • can make use of cloud-based storage for publishing repositories
- 118 • provides useful features aimed specifically at APT repository management
- 119 • allow publishing several repositories and components to a single endpoint

120 Its main shortcoming (local pool) can be addressed by using the REST API for  
121 running `aptly` on a dedicated server. In the future, it might also be possible to  
122 configure a different `aptly` server per OBS project.

123 **Implementation plan**

- 124 • Update OBS to the latest upstream version
- 125 • Start with a prototype, local-only version capable of:
  - 126 – adding a package to a (manually created) local repository
  - 127 – publishing the local repository
  - 128 – deleting a package from the repository when removing it from OBS
- 129 • Implement automated branching and repository creation for new OBS  
130 projects
- 131 • Add configuration options for publishing to cloud-based storage
- 132 • Automate periodic database cleanups
- 133 • Implement REST API interface (global configuration)