



System updates and rollback

1	Contents	
2	Definitions	2
3	Base OS	2
4	Applications	2
5	Use cases	2
6	Embedded device on the field	2
7	Typical system update	2
8	Critical security update	2
9	Applications and base OS with different release cadence	3
10	Shared base OS	3
11	Reselling a device	3
12	Non use cases	3
13	User modified device	3
14	Unrecoverable hardware failure	3
15	Unrecoverable file system corruption	3
16	Development	3
17	Requirements	4
18	Minimal resource consumption	4
19	Work on different hardware platforms	4
20	Every updated system should be identical to the reference	4
21	Atomic update	4
22	Rolling back to the last known good state	4
23	Reset to clean state	5
24	Update control interface	5
25	Existing system update mechanisms	5
26	Debian tools	5
27	ChromeOS	5
28	Approach	6
29	Advantages of using OSTree	6
30	The OSTree model	7
31	Resilient upgrade workflow	9
32	Online web-based OTA updates	10
33	Offline updates	13
34	Online web-based OTA updates using OSTree Static Deltas	14
35	OSTree security	15
36	Error handling	20
37	Implementation	22
38	The general flow	22
39	The boot count	22
40	The bootloader integration	23
41	The updater daemon	24
42	Command line HMI	25
43	Update validation	26
44	Testing	26
45	The update process is robust in case of errors	26

46	Images roll back in case of error	27
47	Images are a suitable rollback target	27
48	User and user data management	27
49	Application management	27
50	Application storage	28
51	Further developments	28
52	Related Documents	29

53 This document focuses on the system update mechanism, but also partly ad-
54 dresses applications and how they interact with it.

55 **Definitions**

56 **Base OS**

57 The core components of the operating system that are used by almost all Apertis
58 users. Hardware control, resource management, service life cycle monitoring,
59 networking

60 **Applications**

61 Components that work on top of the base OS and are specific to certain usages.

62 **Use cases**

63 A variety of use cases for system updates and rollback are given below.

64 **Embedded device on the field**

65 An Apertis device is shipped to a location that cannot be easily accessed by a
66 technician. The device should not require any intervention in the case of errors
67 during the update process and should automatically go back to a know-good
68 state if needed.

69 The update process should be robust against power losses and low voltage situ-
70 ations, loss of connectivity, storage exhaustion, etc.

71 **Typical system update**

72 The user can update his system to run the latest published version of the soft-
73 ware. This can be triggered either via periodic polling, upon user request, or
74 any other suitable mean.

75 **Critical security update**

76 In the case of a critical security issue, the OEM could push an “update avail-
77 able” message to some component in the device that would in turn trigger the
78 update. This requires an infrastructure to reference all devices on the OEM

79 side. The benefit compared to periodic polling is that the delay between the
80 update publication and the update trigger is shortened.

81 **Applications and base OS with different release cadence**

82 Base OS releases involve many moving parts while application releases are sim-
83 pler, so application authors want a faster release cadence decoupled from the
84 base OS one.

85 **Shared base OS**

86 Multiple teams using the same hardware platform want to use the same base
87 OS and differentiate their product purely with applications on top of it.

88 **Reselling a device**

89 Under specific circumstances, the user might want to reset his device to a clean
90 state with no device-specific or personal data. This can happen before reselling
91 the device or the user encountered an unexpected failure.

92 **Non use cases**

93 **User modified device**

94 The user has modified his device. For example, they mounted the file system
95 read write, and tweaked some configuration files to customize some features. As
96 a result, the system update mechanism may no longer be functional.

97 It might still be possible to restore the operating system to a factory state but
98 the system update mechanism cannot guarantee it.

99 **Unrecoverable hardware failure**

100 An hardware failure has damaged the flash storage or another core hardware
101 component and the system is no longer able to boot. Compensating for hardware
102 failures is not part of the system update mechanism.

103 **Unrecoverable file system corruption**

104 The file system became corrupted due to a software bug or other failure and is
105 not able to automatically correct the error. How to recover from that situation
106 is not part of the system update and rollback mechanism.

107 **Development**

108 Developers need to modify and customize their environment in a way that often
109 conflicts with the requirements for devices on the field.

110 **Requirements**

111 **Minimal resource consumption**

112 Some devices only have a very limited amount of available storage, the system
113 update mechanism must keep the impact storage requirement as low as possible
114 and have a negligible impact at runtime.

115 **Work on different hardware platforms**

116 Different devices may use different CPU architectures, bootloaders, storage tech-
117 nologies, partitioning schemas and file system formats.

118 The system update mechanism must be able to work across them with mini-
119 mal changes, ranging from single-partition systems running UBIFS on NAND
120 devices to more common storage devices using traditional file systems over mul-
121 tiple partitions.

122 **Every updated system should be identical to the reference**

123 The file system contents of the base OS on the updated devices must match
124 exactly the file system used during testing to ensure that its behavior can be
125 relied upon.

126 This also means that applications must be kept separate from the base OS to
127 be able to update them while keeping the base OS immutable.

128 **Atomic update**

129 To guarantee robustness in case of errors, every update to the system must be
130 atomic.

131 This means that if an update is not successful, it must not be partially installed.
132 The failure must leave the device in the same state as if the update did not start
133 and no intermediate state must exist.

134 **Rolling back to the last known good state**

135 If the system cannot boot correctly after an update has been installed success-
136 fully it must automatically roll back to a known working version.

137 Applications must be kept separated to be able to roll back the base OS while
138 preserving them or to roll them back while keeping the base OS unchanged.

139 The policy deciding what to roll back and when is product-specific and must
140 be customizable. For instance, some products may chose to only roll back the
141 base OS and keep applications untouched, some other products may choose to
142 roll applications back as well.

143 **Reset to clean state**

144 The user must be able to restore his device to a clean state, destroying all user
145 data and all device-specific system configuration.

146 **Update control interface**

147 An interface must be provided by the updates and rollback mechanism to allow
148 HMI to query the current update status, and trigger updates and rollback.

149 **Existing system update mechanisms**

150 **Debian tools**

151 The Debian package management binds all the software in the system. This can
152 be very convenient and powerful for administration and development, but this
153 level of management is not required for final users of Apertis. For example:

- 154 • Package administration command line tools are not required for final users.
- 155 • No support for update roll back. If there is some package breakage, or
156 broken upgrade, the only way to solve the issue is manually tracking the
157 broken package and downgrading to a previous version, solving dependen-
158 cies along the way. This can be an error prone manual process and might
159 not be accomplished cleanly.

160 **ChromeOS**

161 ChromeOS uses an A/B parallel partition approach. Instead of upgrading the
162 system directly, it installs a fresh image into B partition for kernel and rootfs,
163 then flag those to be booted next time.

164 The partition metadata contains boot fields for the boot attempts (successful
165 boots) and these are updated for every boot. If a predetermined number of
166 unsuccessful boots is reached, the bootloader falls back to the other partition,
167 and it will continue booting from there until the next upgrade is available. When
168 the next upgrade becomes available it will replace the failing installation and
169 will attempt booting from there again.

170 There are some drawbacks to this approach when compared to OSTree:

- 171 • The OS installations are not deduplicated, the system stores the entire
172 contents of the A and B installations separately, where as OSTree based
173 systems only store the base system plus a delta between this and any
174 update using Unix hard links. This means an update to the system only
175 requires disk space proportional to the changed files.
- 176 • The A/B approach can be less efficient since it will need to add extra
177 layers to work with different partitions, for example, using a specific layer
178 to verify integrity of the block devices, where OSTree directly handles

179 operating system views and a content addressable data store (file system
180 user space) avoiding the need of having different layers.
181 • Several partitions are usually required to implement this model, reducing
182 the flexibility with which the storage space in the device can be utilized.


183 Approach

184 Package-based solutions fail to meet the robustness requirements, while dual
185 partitioning schemes have storage requirements that are not viable for smaller
186 devices.

187 OSTree¹ provides atomic updates on top of any POSIX-compatible file system
188 including UBIFS on NAND devices, is not tied to a specific partitioning scheme
189 efficiently handles the available storage.

190 No specific requirements are imposed on the partitioning schema. Use of
191 the GUID Partition Table (GPT²) system for partition management is
192 recommended for being flexible while having fail-safe measures, like keeping
193 checksums of the partition layout and providing some redundancy in case
194 errors are detected.

195 Separating the system volume from the general storage volume, where applica-
196 tions and user data are stored, is also recommended.



197 GPT Partitions

198 More complex schemas can be used for instance by combining OSTree with read-
199 only fallback partitions to handle file system corruption on the main system
200 partition, but this document focuses on a OSTree-only setup that provides a
201 good balance between complexity and robustness.

202 Advantages of using OSTree

- 203 • OSTree operates at the Unix file system layer and thus on top of any
204 file system or block storage layout, including NAND flash setups, and in
205 containers.
- 206 • OSTree does not impose strict requirements on the partitioning scheme
207 and can scale down to a single partition while fully preserving its resiliency

¹<http://ostree.readthedocs.io>

²http://en.wikipedia.org/wiki/GUID_Partition_Table

208 guarantees, saving space on the device and avoiding extra layers of com-
209 plexity (for example, to verify partition blocks). Depending on the setup,
210 multiple partitions can still be used effectively to separate contents with
211 different life cycles, for instance by storing user data on a different parti-
212 tion than the system files managed by OSTree.

- 213 • OSTree commits are centrally created offline (server side), and then they
214 are deployed by the client. This gives much more control over what the
215 devices actually run.
- 216 • It can store multiple file systems trees in a single repository.
- 217 • It is designed to implement fully atomic and resilient upgrades. If the
218 system crashes or power is lost at any point during the update process,
219 you will have either the old system, or the new one.
- 220 • It clearly separate the OS from the device configuration and user data,
221 so resetting the system to a clean state simply involves deleting some
222 directories and their contents.
- 223 • OSTree is implemented as a shared library, making it very easy to build
224 higher level projects or tools on top of it.
- 225 • The files in `/usr` contents are mounted read-only from subfolders of `/os-`
226 `tree/deploy`, minimizing the chance of accidental deletions or changes.
- 227 • OSTree has no impact on startup performance, nor does increase resource
228 usage during runtime: since OSTree is just a different way to build the
229 rootfs once it is built it will behave like a normal rootfs, making it very
230 suitable for setups with limited storage.
- 231 • OSTree already offers a mechanism suitable for offline updates using static
232 deltas, which can be used for updates via a mass-storage device.
- 233 • Security is at the core of OSTree, offering content replication incrementally
234 over HTTPS via GPG signatures and using SHA256 hash checksums.
- 235 • The mechanism to apply partial updates or full updates is exactly the
236 same, the only difference is how the updates are generated on the server
237 side.
- 238 • OSTree can be used for both the base OS and applications, and its built-in
239 hard link-based deduplication mechanism allow to share identical contents
240 between the two, to keep them independent while having minimal impact
241 on the needed storage. The Flatpak application framework is already
242 based on OSTree.

243 **The OSTree model**

244 The conceptual model behind OSTree repositories is very similar to the one used
245 by `git`, to the point that the [introduction in the OSTree manual](#)³ refers to it as
246 “git for operating system binaries”.

247 Albeit they take different tradeoffs to address different use-cases they both have:

- 248 • file contents stored as blobs addressable by their hash, deduplicating them

³<https://ostree.readthedocs.io/en/latest/manual/introduction/>

- 249 • file trees linking filenames to the blobs
- 250 • commits adding metadata such as dates and comments on top of file trees
- 251 • commits linked in a history tree
- 252 • branches pointing to the most recent commit in a history tree, so that
- 253 clients can find them

254 Where `git` focuses on small text files, OSTree focuses on large trees of binary
255 files.

256 On top of that OSTree adds other layers which go beyond storing and distribut-
257 ing file contents to fully handle operating system upgrades:

- 258 • repositories - store one or more versions of the file system contents as
- 259 described above
- 260 • deployments - specific file system versions checked-out from the repository
- 261 • stateroots - the combination of immutable deployments and writable di-
- 262 rectories

263 Each device hosts a local OSTree repository with one or more deployments
264 checked out.

265 Checked out deployments look like traditional root file systems. The bootloader
266 points to the kernel and initramfs carried by the deployment which, after setting
267 up the writable directories from the stateroot, are responsible for booting the
268 system. The bootloader is not part of the updates and remains unchanged for
269 the whole lifetime of the device as any changes has a high chance to make the
270 system unbootable.

- 271 • Each deployment is grouped in exactly one `stateroot`, and in normal cir-
272 cumstances Apertis devices only have a single `apertis` `stateroot`.
- 273 • A `stateroot` is physically represented in the `/ostree/deploy/$stateroot` di-
274 rectory, `/ostree/deploy/apertis` in this case.
- 275 • Each `stateroot` has exactly one copy of the traditional Unix `/var` directory,
276 stored physically in `/ostree/deploy/$stateroot/var`. The `/var` directory is
277 persisted during updates, when moving from one deployment to another
278 and it is up to each operating system to manage this directory.
- 279 • On each device there is an OSTree repository stored in `/ostree/repo`, and
280 a set of deployments stored in `/ostree/deploy/$stateroot/$checksum`.
- 281 • A deployment begins with a specific `commit` (represented by a SHA256
282 hash) in the OSTree repository in `/ostree/repo`. This `commit` refers to a file
283 system tree that represents the underlying basis of a deployment.
- 284 • Each deployment is primarily composed of a set of hard links into the
285 repository. This means each version is deduplicated; an upgrade process
286 only costs disk space proportional to the new files, plus some constant
287 overhead.
- 288 • The read-only base OS contents are checked out in the `/usr` directory of
289 the deployment.
- 290 • Each deployment has its own writable copy of the configuration store `/etc`.

- 291 • Deployments don't have a traditional UNIX `/etc` but ship it instead as
292 `/usr/etc`. When OSTree checks out a deployment it performs a 3-way
293 merge using the old default configuration, the active system's `/etc`, and
294 the new default configuration.
- 295 • Besides the exceptions of `/var` and `/etc` directories, the rest of the contents
296 of the tree are checked out as hard links into the repository.
- 297 • Both `/etc` and `/var` are persistent writable directories that get preserved
298 across upgrades.

299 Resilient upgrade workflow

300 The following steps are performed to upgrade a system using OSTree:

- 301 • The system boots using the existing deployment
- 302 • A new version is made available as a new OSTree commit in the local
303 repository, either downloading it from the network or by unpacking a
304 static delta shipped on a mass storage device.
- 305 • The data is validated for integrity and appropriateness.
- 306 • The new version is deployed.
- 307 • The system reboots into the new deployment.
- 308 • If the system fails to boot properly (which should be determined by the
309 system boot logic), the system can roll back to the previous deployment.

310 During the upgrade process, OSTree will take care of many important details,
311 like for example, managing the bootloader configuration and correctly merging
312 the `/etc` directory.

313 Each `commit` can be delivered to the target system over the air or by attaching a
314 mass storage device. Network upgrades and mass storage upgrades only differ
315 in the mechanism used by `ostree` to detect and obtain the update. In both cases
316 the `commit` is first stored in a temporary directory, validated and only then it
317 becomes part of the local OSTree repository before the real upgrade process
318 starts by rebooting in the new deployment.

319 Metadata such as EdDSA or GPG signatures can be attached to each `commit` to
320 validate it, ensuring it is appropriate for the current system and it has not been
321 corrupted or tampered. The update process must be interrupted at any point
322 during the update process should any check yield an invalid result; the [atomic
323 upgrades mechanism in OSTree](#)⁴ ensures that it is safe to stop the process at
324 any point and no change is applied to the system up to the last step in the
325 process.

326 The atomic upgrades mechanism in OSTree ensures that any power failure dur-
327 ing the update process leaves the current system state unchanged and the up-
328 date process can be resumed re-using all the data that has already been already
329 validated and included in the local repository.

⁴<https://ostree.readthedocs.io/en/latest/manual/atomic-upgrades/#atomic-upgrades>

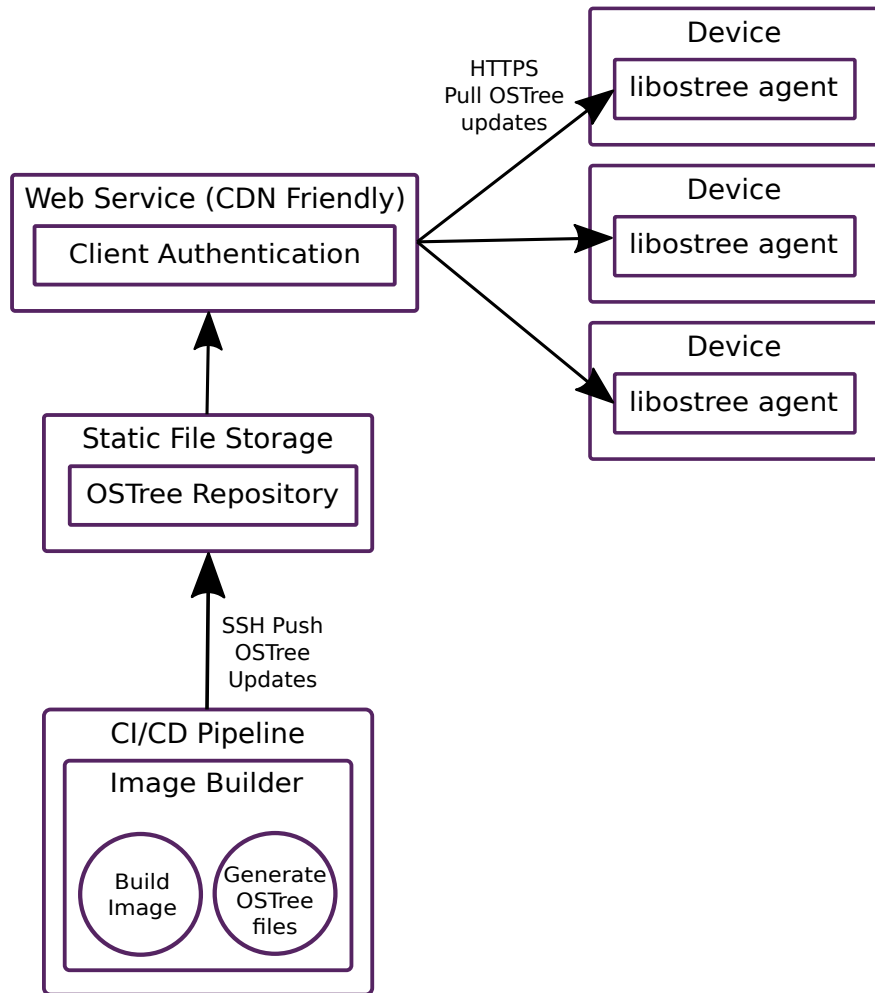
330 **Online web-based OTA updates**

331 OSTree supports bandwidth-efficient retrieval of updates over the network.

332 The basic workflow involves the actors below:

- 333 • the image building pipelines pushes commits to an OSTree repository on
334 each build;
- 335 • a standard web server provides access over HTTPS to the OSTree reposi-
336 tory handling it as a plain hierarchy of static files, with no special knowl-
337 edge of OSTree;
- 338 • the client devices poll the web server and retrieve updates when they get
339 published.

340 The following diagram shows how the data flows across services when using the
341 web based OSTree upgrade mechanism.



342

343 Thanks to its repository format, OSTree client devices can efficiently query the
 344 repository status and retrieve only the changed contents without any OSTree-
 345 specific support in the web server, with the repository files being served as plain
 346 static files.

347 This means that any web hosting provider can be used without any loss of
 348 efficiency.


349 By only requiring static files, the web service can easily take advantage of CDN
 350 services to offer a cost efficient solution to get the data out to the devices in a
 351 way that is globally scalable.

352 The authentication to the web service can be done via HTTP Basic authentica-

353 tion, SSL/TLS client certificates, or any cookie-based mechanism that is most
354 suitable for the chosen web service, as OSTree does not impose any constraint
355 over plain HTTPS. OSTree separately checks the chain of trust linking the down-
356 loaded updates to the keys trusted by the system update manager. See also the
357 [Controlling access to the updates repository](#) and [Verified updates](#) sections in
358 this regard.

359 Monitoring and management of devices can be built using the same HTTPS
360 access as used by OSTree or using completely separated mechanisms, enabling
361 the integration of OSTree updates into existing setups.

362 For instance, integration with roll out management suites like [Eclipse hawkBit](#)⁵
363 can happen by disabling the polling in the OSTree updater and letting the man-
364 agement suite tell OSTree which commit to download and from where through
365 a dedicated agent running on the devices.



366 This has the advantage that the roll out management suite would be in complete
367 control of which updates should be applied to which devices, implementing
368 any kind of policies like progressive staged roll outs with monitoring from the
369 backend with minimal integration.
370

371 Only the retrieval and application of the actual update data on the device
372 would be offloaded to the OSTree-based update system, preserving its network
373 and storage efficiency and the atomicity guarantees.

⁵<https://www.eclipse.org/hawkbit/>

374 **Offline updates**

375 Some devices may not have any connectivity, or bandwidth requirements may
376 make full system updates prohibitive. In these cases updates can be made avail-
377 able offline by providing OSTree “static delta” files on external media devices
378 like USB mass storage devices.

379 The deltas are simple static files that contains all the differences between two
380 specific OSTree commits. The user would download the delta file from a web
381 site and put it in the root of an external drive. After the drive is mounted, the
382 update management system would look for files with a specific name pattern in
383 the root of the drive. If an appropriate file is found, it is checked to be a valid
384 OSTree static bundle with the right metadata and, if that verification passes, the
385 user would get a notification saying that updates are available from the drive.
386 If the update file is corrupted, is targeted to other platforms or devices, or is
387 otherwise invalid, the upgrade process must stop, leaving the system unchanged
388 and a notification may be reported to the user about the identified issues.

389 Static deltas can be partial if the devices are known beforehand to have a specific
390 OSTree commit already available locally, or they can be full by providing the
391 delta from the `NULL` empty commit, thus ensuring that the update can be applied
392 without any assumption on the version running on the devices at the expense
393 of a potential increase of the requirements on the mass storage device used to
394 ship them. Both partial and full deltas leading to the same OSTree commit will
395 produce identical results on the devices.

396 **Switching to the new branch**

397 The branches naming schema used in Apertis contains the major version, for in-
398 stance: `apertis/v2020/armhf-uboot/minimal`. So for Apertis the “major upgrade”
399 is technically considered as switching to another branch with a more recent
400 Apertis version, for example `apertis/v2021/armhf-uboot/minimal`. By default such
401 kinds of offline upgrade with switching to another branch is restricted by the
402 update manager.

403 Offline upgrades between branches (including “major updates”) consists of 2
404 steps which should be a part of offline upgrade:

- 405 1. Prepare the proper commit at build time

406 This step is pretty simple – while preparing the relevant commit we just
407 need to add the branch name(s) from which we are supposed to be able
408 to upgrade to the current version. For example, while preparing commit
409 for `v2021` version just add following into `ostree-commit` action:

```
1 ref-binding:  
2   - apertis/v2021/{{ $architecture }}-{{ $board }}/{{ $type }}  
3   - apertis/v2020/{{ $architecture }}-{{ $board }}/{{ $type }}
```

410 this produce the commit compatible with version `v2020` allowing to install
411 the new OS version `v2021` onto the board.

412 2. Set correct refs in repository

413 After successful boot of the updated version all refs in `libostree` repository
414 are still pointing to the previous branch due the nature of offline upgrade.

415 This needs to be fixed for proper detection of further upgrades, including
416 updates over the air. It is the responsibility of the update manager to
417 update the refs once the update has been determined to be successful.

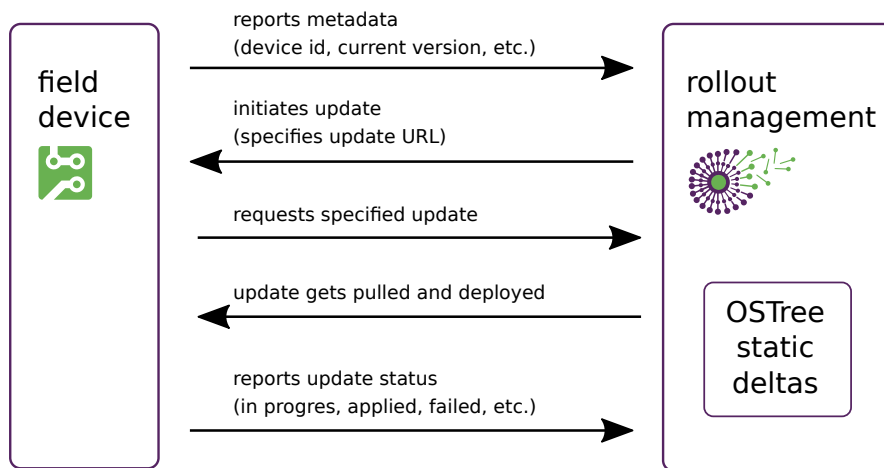
418 This functionality requires no changes to be made to previously released OSTree
419 versions. The configuration that determines upgrade paths is held in the newer
420 OSTree commit.

421 Apertis currently only supports upgrades to newer versions, downgrades to older
422 versions of Apertis are not supported.

423 Online web-based OTA updates using OSTree Static Deltas

424 The OSTree based OTA update mechanism described above, whilst taking full
425 advantage of OSTree's repositories, may not suit all users of Apertis. Where
426 existing deployment services such as hawkBit are already deployed, a need to
427 expose and maintain an extra service, the OSTree repository, may not be
428 welcome.

429 The ability of OSTree to perform updates using static deltas provides us with
430 the option to serve these via the deployment infrastructure instead of serving
431 updates from the OSTree repository.



432

433 Such an approach would likely need to utilize full OSTree deltas in order to
434 ensure updates were viable on target devices regardless of how frequently they

435 had previously been updated. As a result this approach does not take advantage
436 of the bandwidth efficiency that would be presented by **Online web-based OTA**
437 **updates** as previously discussed, however it does enable updates to be performed
438 via existing installed deployment infrastructure.

439 This is the approach currently used by the Apertis example implementation.

440 **OSTree security**

441 OSTree is a distribution method. It can secure the downloading of the update
442 by verifying that it is properly signed using public key cryptography (EdDSA
443 or GPG). It is largely orthogonal to verified boot, that is ensuring that only
444 signed data is executed by the system from the bootloader, to the kernel and
445 user space. The only interaction is that since OSTree is a file-based distribution
446 mechanism, block-based verification mechanism like `dm-verity` cannot be used.
447 OSTree can be used in conjunction with signed bootloader, signed kernel, and
448 IMA (Integrity Measurement Architecture) to provide protection from offline
449 attacks.

450 **Verified boot**

451 Verified boot is the process which ensures a device is only runs signed code.
452 This is a layered process where each layer verifies signature of its upper layer.

453 The bottom layer is the hardware, which contains a data area reserved to certifi-
454 cates. The first step is thus to provide a signed bootloader. The processor can
455 verify the signature of the bootloader before starting it. The bootloader then
456 reads the boot configuration file. It can then run a signed kernel and `initramfs`.
457 Once the kernel has started, the `initramfs` mounts the root file system.

458 At the time of writing, the boot configuration file is not signed. It is read and
459 verified by signed code, and can only point to signed components.

460 Protecting bootloader, kernel and `initramfs` already guarantees that policies
461 baked in those components cannot be subverted through offline attacks. By
462 verifying the content of the rootfs the protection can be extended to user space
463 components, albeit such protection can only be partial since device-local data
464 can't be signed on the server-side like the rest of the rootfs.

465 To protect the rootfs different mechanisms are available: the block-based ones
466 like `dm-verity` are fundamentally incompatible with file-based distribution meth-
467 ods like OSTree, since they rely on the kernel to verify the signature on each
468 read at the block level, guaranteeing that the whole block device has not been
469 changed compared to the version signed at deployment time. Due to working
470 on raw block devices, `dm-verity` is also incompatible with UBIFS and thus it is
471 unsuitable for NAND devices.

472 Other mechanisms like IMA (Integrity Measurement Architecture) work instead
473 at the file level, and thus can be used in conjunction with UBIFS and OSTree

474 on NAND devices.

475 It is also possible to check that the deployed OSTree rootfs matches the server-
476 provided signature without using any other mechanism, but unlike IMA and
477 `dm-verity` such check would be too expensive to be done during file access.

478 **Verified updates**

479 Once a verified system is running, an OSTree update can be triggered. Apertis is
480 using [ed25519](#)⁶ variant of EdDSA signature. Ed25519 ensures that the commit
481 was not modified, damaged, or corrupted.

482 On the server, OSTree commits must be signed using ed25519 secret key. This
483 occurs via the `ostree sign --sign-type=ed25519 <COMMIT_ID>` command line. The
484 secret key could be provided via additional CLI parameter or file by using option
485 `--keys-file=<path_to_file>`.

486 OSTree expect what secret key consists of 64 bytes (32b seed + 32b public)
487 encoded with base64 format. The ed25519 secret and public parts could be
488 generated by numerous utilities including `openssl`, for instance:

```
1  openssl genpkey -algorithm ed25519 -outform PEM -out ed25519.pem
```

489 Since OSTree is not capable to use PEM format directly, it is needed to [extract](#)
490 [the secret and public keys](#)⁷ from PEM file, for example:

```
1  PUBLIC="$(openssl pkey -outform DER -pubout -in ${PEMFILE} | tail -  
2  c 32 | base64) "  
   SEED="$(openssl pkey -outform DER -in ${PEMFILE} | tail -c 32 | base64) "
```

491 As mentioned above, the secret key is concatenation of SEED and PUBLIC
492 parts:

```
1  SECRET="$(echo ${SEED}${PUBLIC} | base64 -d | base64 -w 0) "
```

493 On the client, ed25519 is also used to ensure that the commit comes from a
494 trusted provider since updates could be acquired through different methods like
495 OTA over a network connection, offline updates on plug-in mass storage devices,
496 or even mesh-based distribution mechanism. To enable the signature check,

⁶<https://ed25519.cr.yp.to/>

⁷<http://openssl.6102.n7.nabble.com/ed25519-key-generation-td73907.html>

497 repository on the client must be configured by adding option `sign-verify=true`
498 into the `core` or `per-remote` section, for instance:

```
1  ostree config set 'remote "origin".sign-verify' "true"
```

499 OSTree searches for files with valid public signatures in directories
500 `/usr/share/ostree/trusted.ed25519.d` and `/etc/ostree/trusted.ed25519.d`.
501 Any public key in a file in these directories will be trusted by the client. Each
502 file may contain multiple keys, one base64-encoded public key per string. No
503 private keys should be present in these directories.

504 In addition it is possible to provide the trusted public key per-remote by
505 adding into remote's configuration path to the file with trusted public keys (via
506 `verification-file` option) or even single key itself (via `verification-key`).

507 In the OSTree configuration, the default is to require commits to be signed.
508 However, if no public key is available, no any commit can be trusted.

509 **Offline update files with signed metadata**

510 Starting with version v2020.7 `libostree` supports delta bundles with [signed meta-](#)
511 [data](#)⁸, which allows to ensure that the whole delta bundle comes from a trusted
512 source. Previous versions only allowed to assert the provenance of the commits
513 in the bundle, leaving the metadata unverified.

514 Support for delta bundles with signed metadata is available in the Apertis Up-
515 date Manager since version 0.2020.20, which can also handle delta bundles with
516 unsigned metadata. Previous versions of the Apertis Update Manager also
517 supported an experimental format for signed metadata, which has now been
518 dropped in favor of the format that has been landed upstream in `libostree`
519 2020.7.

520 To improve the security on target devices the repository configuration must have
521 additional option `core.sign-verify-deltas` set to `true`:

```
1  ostree config set core.sign-verify-deltas "true"
```

522 This is forcing AUM and `libostree` to accept only update bundles with signed
523 metadata.

524 **Compatibility with upgrades**

⁸<https://github.com/ostreedev/ostree/pull/1985>

525 Until `v2021pre` there were no support of upgrade bundles with signed metadata
526 in Apertis.

527 For upgrading from version `v2020` to `v2021` we have produced additional upgrade
528 bundle. This additional bundle has unsigned metadata, allowing offline upgrades
529 from the previous release. So for `v2021`, and only for `v2021` we have 3 update
530 bundle types:

- 531 • `*.delta` – depending on CI it may have signed or unsigned metadata. This
532 version is uploaded into hawkBit server and used for tests.
- 533 • `*.delta.enc` – encrypted bundle containing delta file above.
- 534 • `*.compat-v2020.delta` – bundle with unsigned metadata compatible with
535 previous Apertis release. This file should be used for upgrading the Apertis
536 version `v2020`.

537 **Securing OSTree updates download**

538 OSTree supports “pinned TLS”. Pinning consist of storing the public key of the
539 trusted host on the client side, thus eliminating the need for a trust authority.

540 TLS can be configured in the `remote` configuration on the client using the follow-
541 ing entries:

```
1  tls-ca-path
2      Path to file containing trusted anchors instead of the system CA database.
```

542 Once a key is pinned, OSTree is ensured that any download is coming from a
543 host which key is present in the image.

544 The pinned key can be provided in the disk image, ensuring every flashed device
545 is able to authenticate updates.

546 **Controlling access to the updates repository**

547 TLS also permit the OSTree client to authenticate itself to the server before
548 being allowed to download a commit. This can also be configured in the `remote`
549 configuration on the client using the following entries:

```
1  tls-client-cert-path
2          Path to file for client-
3  side certificate, to present when making requests to
4  this repository.
5  tls-client-key-path
6          Path to file containing client-
   side certificate key, to present when making
   requests to this repository.
```

550 Access to remote repositories can also be controlled via HTTP cookies. The
551 `ostree remote add-cookie` and `ostree remote delete-cookie` commands will up-
552 date a per-remote lookaside cookie jar, named `$remotename.cookies.txt`. In this
553 model, the client first obtains an authentication cookie before communicating
554 this cookie to the server along with its update request.

555 The choice between authentication via TLS client-side certificates or HTTP
556 cookies can be done depending on the chosen server-side infrastructure.

557 Provisioning authentication keys on a per-device basis at the end of the deliv-
558 ery chain is recommended so each device can be identified and access granted
559 or denied at the device granularity. Alternatively it is possible to deploy au-
560 thentication keys at coarser granularities, for instance one for each device class,
561 depending on the specific use-case.

562 **Security concerns for offline updates over external media**

563 OSTree static deltas includes the detached metadata with signature for the
564 contained commit to check if the commit is provided by a valid provider and its
565 integrity.

566 The signed commit is unpacked to a temporary directory and verified by OSTree
567 before being integrated in the OSTree repository on the device, from which it
568 can be deployed at the next reboot.

569 This is the same mechanism used for commit verification when doing OTA
570 upgrades from remote servers and provides the same features and guarantees.

571 Usage of inlined signed metadata ensures that the provided update file is aimed
572 to the target platform or device.

573 Updates from external media present a security problem not present for directly
574 downloaded updates. Simply verifying the signature of a file before decompress-
575 ing is an incomplete solution since a user with sufficient skill and resources
576 could create a malicious USB mass storage device that presents different data
577 during the first and second read of a single file – passing the signature test, then
578 presenting a different image for decompression.

579 The content of the update file is extracted into the temporary directory and the
580 signature is checked for the extracted commit tree.

581 **Error handling**

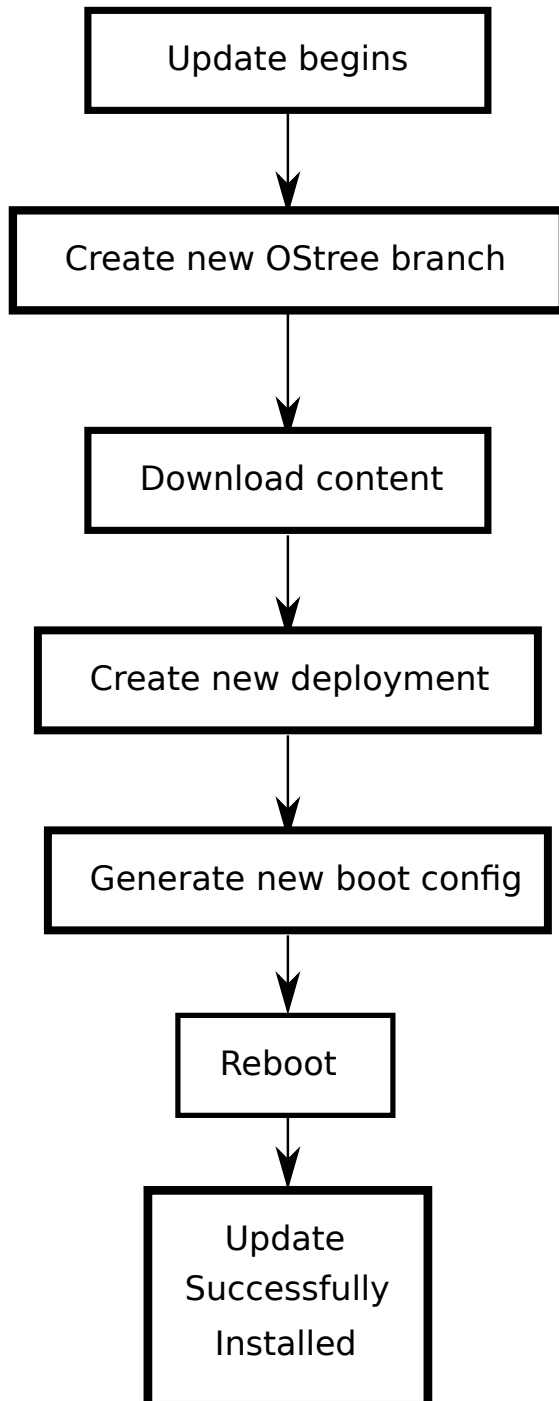
582 If for any reason the update process fails to complete, the update will be black-
583 listed to avoid re-attempting it. Another update won't be automatically at-
584 tempted until a newer update is made available.

585 The only exception from this rule is failure due incorrect signature check. The
586 commit could be re-signed with the key not known for the client at this moment,
587 and as soon as client acquire the new public key blacklist mechanism shouldn't
588 prevent the update.

589 It is possible that an update is successfully installed yet fail to boot, resulting in
590 a rollback. In the event of a rollback the update manager must detect that the
591 new update has not been correctly booted, and blacklist the update so it is not
592 attempted again. To detect a failed boot a watchdog mechanism can be used.
593 The failed updates can then be blacklisted by appending their OSTree commit
594 ids to a list.

595 This policy prevents a device from getting caught in a loop of rollbacks and
596 failed updates at the expense of running an old version of the system until a
597 newer update is pushed.

598 The most convenient storage location for the blacklist is the user storage area,
599 since it can be written at runtime. As a side effect of storing the blacklist there,
600 it will automatically be cleared if the system is reset to a clean state.



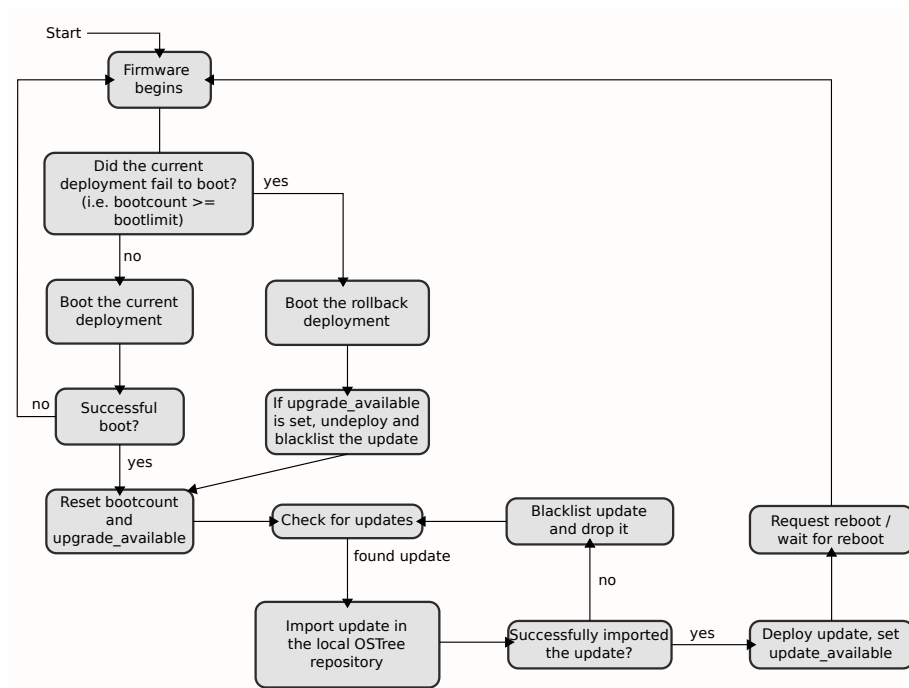
602 Implementation

603 This section provides some more details about the implementation of offline
604 system updates and rollback in Apertis, which is split in three main components:

- 605 • the updater daemon
- 606 • the bootloader integration
- 607 • the command-line HMI

608 The general flow

609 The Apertis update process deals with selecting the OSTree deployment to boot,
610 rolling back to known-good deployments in case of failure and preparing the new
611 deployment on updates:



612
613 While the underlying approach differs due to the use of OSTree in Apertis over
614 the dual-partition approach chosen by ChromeOS and the different bootloaders,
615 the update/rollback process is largely the same as [the one in ChromeOS⁹](https://www.chromium.org/chromium-os/chromiumos-design-docs/filesystem-autoupdate#TOC-Diagram).

616 The boot count

617 To keep track of failed updates the system maintains a persistent counter that
618 it is increased every time a boot is attempted.

⁹<https://www.chromium.org/chromium-os/chromiumos-design-docs/filesystem-autoupdate#TOC-Diagram>

619 Once a boot is considered successful depending on project-specific policies (for
620 instance, when a specific set of services has been started with no errors) the
621 boot count is reset to zero.

622 This boot counter needs to be handled in two places:

- 623 • in the bootloader, which boots the current OSTree deployment if the
624 counter is zero and initiates a rollback otherwise
- 625 • in the updater, which needs to reset it to zero once the boot is considered
626 successful

627 Using the main persistent storage to store the boot count is viable for most
628 platform but would produce too much wear on platforms using NAND devices.
629 In those cases the boot count should be stored on another platform-specific
630 location which is persistent over warm reboots, there's no need for it to persist
631 over power losses.

632 However, in the reference implementation the focus is on the most general solu-
633 tion first, while being flexible enough to accommodate other solutions whenever
634 needed.

635 **The bootloader integration**

636 Since bootloaders are largely platform-specific the integration needs to be done
637 per-platform.

638 For the SabreLite ARM 32bit platform, integration with the [U-Boot](#)¹⁰ boot-
639 loader is needed.

640 OSTree already provides dedicated hooks to update the `u-boot` environment to
641 point it to the latest deployment.

642 Two separate boot commands are used to start the system: the default one boots
643 the latest deployment, while the alternate one boots the previous deployment.

644 Before rebooting to a new deployment the boot configuration file is switched
645 and the one for the new deployment is made the default, while the older one is
646 put into the location pointed by the alternate boot command.

647 When a failure is detected by checking the boot count while booting the latest
648 deployment, the system reboots using the alternate boot command into the
649 previous deployment where the rollback is completed.

650 Once the boot procedure completes successfully the boot count gets reset and
651 stopped, so failures in subsequent boots won't cause rollbacks which may worsen
652 the failure.

653 If the system detects that a rollback has been requested, it also need to make
654 the rollback persistent and prevent the faulty updates to be tried again. To do

¹⁰<http://www.denx.de/wiki/U-Boot/WebHome>

655 so, it adds any deployment more recent than the current one to a local blacklist
656 and then drops them.

657 **The updater daemon**

658 The updater daemon is responsible for most of the activities involved, such as
659 detecting available updates, initiating the update process and managing the
660 boot count.

661 It handles both online OTA updates and offline updates made available on
662 locally mounted devices.

663 **Detecting new available updates**

664 For offline updates, the [GVolumeMonitor](https://developer.gnome.org/gio/stable/GVolumeMonitor.html)¹¹ API provided by GLib/GIO is used
665 to detect when a mass storage device is plugged into the device, and the [GFile](https://developer.gnome.org/gio/stable/GFile.html)¹²
666 GLib/GIO API is used to scan for the offline update stored as a plain file in the
667 root of the plugged file system named `static-update.bundle`.

668 For online OTA updates, the [OstreeSysrootUpgrader](https://github.com/ostreedev/ostree/blob/master/src/libostree/ostree-sysroot-upgrader.c)¹³ is used to poll the remote
669 repository for new commits in the configured branch.

670 When combined with roll out management systems like [Eclipse hawkBit](https://www.eclipse.org/hawkbit/)¹⁴, the
671 roll out management agent on the device will initiate the upgrade process with-
672 out the need for polling.

673 **Initiating the update process**

674 Once the update is detected, it is verified and compared against a local blacklist
675 to skip updates that have already failed in the past (see [Update validation]).

676 In the offline case the static delta file is checked for consistency before being
677 unpacked in the local OSTree repository.

678 During online updates, files are verified as they get downloaded.

679 In both cases the new update results in a commit in the local OSTree repository
680 and from that point the process is identical: a new deployment is created from
681 the new commit and the bootloader configuration is updated to point to the
682 new deployment on the next boot.

683 **Reporting the status to interested clients**

684 The updater daemon exports a simple D-Bus interface which allows to check
685 the state of the update process and to mark the current boot as successful.

¹¹<https://developer.gnome.org/gio/stable/GVolumeMonitor.html>

¹²<https://developer.gnome.org/gio/stable/GFile.html>

¹³<https://github.com/ostreedev/ostree/blob/master/src/libostree/ostree-sysroot-upgrader.c>

¹⁴<https://www.eclipse.org/hawkbit/>

686 **Resetting the boot count**

687 During the boot process the boot count is reset to zero using an interface that
688 abstracts over the platform-specific approach.

689 While product-specific policies dictate when the boot should be considered
690 successful, the reference images consider a boot to be successful if the `multi-`
691 `user.target` target is reached.

692 **Marking deployments**

693 Rolled back deployments are added to a blacklist to avoid trying them again
694 over and over.

695 Deployments that have booted successfully get marked as known good so that
696 they are never rolled back, even if at a later point a failure in the boot process is
697 detected. This is to avoid transient issues causing an unwanted rollback which
698 may make the situation worse.

699 To do so, the boot counting is stopped once the current boot is considered
700 successful, effectively marking the current boot as known-good without the need
701 to maintain a whitelist and synchronize it with the bootloader.

702 As a part of marking the deployment as successful the updater daemon checks
703 the target branches from the OSTree commit metadata. If the booted deploy-
704 ment contains references to several branches, the updater daemon determines
705 which branch has the highest version and resets all refs and the origin to that
706 branch. Using the branches naming scheme used in Apertis, this could be con-
707 sidered as “major upgrade” of the system. From this point on, the system is
708 fully switched to the new branch and accepts only upgrades created in the new
709 branch.

710 **Command line HMI**

711 A command line tool is provided to query the status using [the `org.apertis.ApertisUpdateManager`](#)
712 [D-Bus API](#)¹⁵:

```
1  $ updatectl  
2  ** Message: Network connected: No  
3  ** Message: Upgrade status: Deploying
```

713 The current API exposes information about whether the updater is idle, an
714 update is being checked, retrieved or deployed, or whether a reboot is pending
715 to switch to the updated system.

¹⁵<https://gitlab.apertis.org/appfw/apertis-update-manager/blob/master/data/apertis-update-manager-dbus.xml>

716 It can also be used to mark the boot as successful:

```
1 $ updatectl --mark-update-successful
```

717 **Update validation**

718 Before installing updates the updater check their validity and appropriateness
719 for the current system, using the metadata carried by the update itself as pro-
720 duced by the build pipeline. It ensures that the update is appropriate for the
721 system by verifying that the collection id in the update matches the one config-
722 ured for the system. This prevents installing an update meant for a different
723 kind of device, or mixing variants. The updater also checks that the update ver-
724 sion is newer than the one on the system, to prevent downgrade attacks where
725 a older update with known vulnerabilities is used to gain privileged access to a
726 target.

727 **Testing**

728 Testing ensures that the following system properties for each image are main-
729 tained:

- 730 • the image can be updated if a newer update bundle is plugged in
- 731 • the update process is robust in case of errors
- 732 • the image initiates a rollback to a previous deployment if an error is de-
733 tected on boot
- 734 • the image can complete a rollback initiated from a later deployment

735 To do so, a few components are needed:

- 736 • system update bundles have to be built as part of the daily build pipeline
- 737 • a know-good test update bundle with a very large version number must
738 be create to test that the system can update to it

739 At least initially, testing is done manually. Automation from LAVA will be
740 researched later.

741 **Images can be updated**

742 Plugging a device with the known-good test update on it bundle the expectation
743 is that the system detects it, initiates the update and on reboot the deployment
744 from the known-good test bundle is used.

745 **The update process is robust in case of errors**

746 To test that errors during the update process don't affect the system, the device
747 is unplugged while the update is in progress. Re-plugging it after that checks
748 that updates are gracefully restarted after transient errors.

749 **Images roll back in case of error**

750 Injecting an error in the boot process checks that the image initiates the roll
751 back to a previous deployment. Since a newly-flashed image doesn't have any
752 previous deployment available, one needs to be manually set up beforehand by
753 downloading an older OSTree commit.

754 **Images are a suitable rollback target**

755 A known-bad deployment similar to the known-good one can be used to ensure
756 that the current image works as intended when it is the destination of a rollback
757 initiated by another deployment.

758 After updating to the known-bad deployment the system should rollback to the
759 version under test, which should then complete the rollback by cleaning the
760 boot count, blacklisting the failed update and undeploy it.

761 **User and user data management**

762 As described in the [Multiuser](#)¹⁶ design document, Apertis is meant to accom-
763 modate multiple users on the same device, using existing concepts and features
764 of the several open source components that are being adopted.

765 All user data should be kept in the general storage volume on setups where it is
766 available, as it enables simpler separation of concerns, and a simpler implemen-
767 tation of user data removal.

768 Rolling back user and application data cannot be generally applied and no
769 existing general purpose system supports it. Applications must be prepared to
770 handle configuration and data files coming from later versions and handle that
771 gracefully, either ignoring unknown parameter or falling back to the default
772 configuration if the provided one is not usable.

773 Specific products can choose to hook to the update system and manage their
774 own data rollback policies.

775 **Application management**

776 Application management on Apertis has requirements that the main update
777 management system does not:

- 778 • It is unreasonable to expect a system restart after an application update.
- 779 • Each application must be tracked independently for rollbacks. System
780 updates only track one “stream” of rollbacks, where the application update
781 framework must track many.

782 Flatpak matches the requirements and is also based on OSTree. The ability
783 to deduplicate contents between the system repository and the applications

¹⁶<https://sjoerd.pages.apertis.org/apertis-website/concepts/multiuser/>

784 decouples applications from the base OS yet keeping the impact on storage
785 consumption minimal.

786 **Application storage**

787 Applications can be stored per-device or per-user depending on the needs of the
788 product.

789 An application may require storage space for personal settings, license informa-
790 tion, caches, and any manner of long term private storage. These files should
791 generally not be easily accessible to the user as directly modifying them could
792 have detrimental effects on the application.

793 Application storage requirements can be divided into broad groups:

- 794 • An area for application exports to integrate with the system. This is
795 managed by the application manager and not directly by applications
796 themselves.
- 797 • User specific application data – for settings and any other per-user files.
798 In the event of an application rollback, depending on the product this data
799 may get rolled back with the application or the application needs to deal
800 with potentially mismatching versions.
- 801 • Application specific application data – for data that is rolled back with
802 an application but isn't tied to a user account – such as voice samples or
803 map data. This data should be handled in the same way as user specific
804 application data.
- 805 • Cache – easily recreated data. To save space, this should not be stored for
806 rollback purposes, and should be cleared on a rollback in case applications
807 change their cache data formats between versions.
- 808 • Storage for files in standard formats that aren't tied to specific applica-
809 tions, as explained in the [Multiuser](#)¹⁷ design, this storage is shared between
810 all users. This data should be exempt from the rollback system.

811 **Further developments**

- 812 • Handling a larger threat model using [The Update Framework Specifica-](#)
813 [tion](#)¹⁸ / [Uptane](#)¹⁹ with [Aktualizr](#)²⁰
- 814 • Integrating with server side management services like [Eclipse hawkBit](#)²¹
- 815 • Hardware-assisted [verified boot](#)²² with TPM/OP-TEE

¹⁷<https://sjoerd.pages.apertis.org/apertis-website/concepts/multiuser/>

¹⁸<https://github.com/theupdateframework/specification/blob/master/tuf-spec.md>

¹⁹<https://uptane.github.io/>

²⁰<https://foundries.io/insights/2018/05/25/ota-part-1/>

²¹<https://www.eclipse.org/hawkbit/>

²²<https://www.chromium.org/chromium-os/chromiumos-design-docs/verified-boot>

- 816 • File system-level integrity checks [Integrity Measurement Architecture](#)
817 [\(IMA\)/Extended Verification Module \(EVM\)](#)²³
818 • Add fail safe partition to handle file system corruption

819 **Related Documents**

820 A survey of system update managers:

- 821 • https://wiki.yoctoproject.org/wiki/System_Update

822 The OSTree bootable file systems tree store:

- 823 • <http://ostree.readthedocs.io>

824 The U-Boot Bootloader:

- 825 • <http://www.denx.de/wiki/U-Boot/WebHome>

826 The ChromeOS auto-update system:

- 827 • [https://www.chromium.org/chromium-os/chromiumos-design-docs/filesystem-](https://www.chromium.org/chromium-os/chromiumos-design-docs/filesystem-autoupdate)
828 [autoupdate](https://www.chromium.org/chromium-os/chromiumos-design-docs/filesystem-autoupdate)

²³<https://sourceforge.net/p/linux-ima/wiki/Home/>