



Robustness

1	Contents	
2	Introduction	2
3	Requirements	2
4	Approach	2
5	Application data	2
6	General guidelines	3
7	SQLite	3
8	Tracker	4
9	User settings	4
10	Media	4
11	Caches	4
12	Filesystems	5
13	Root filesystem	8
14	Other filesystems	8
15	Main storage	8
16	Removable devices	9
17	Mitigating the effects of lack of disk space	10
18	Resource management	10
19	CPU	11
20	I/O	12
21	Memory	12
22	Network queue	13
23	GPU	14
24	Accounting	14
25	USB undervoltage	14
26	Risks	14
27	Design notes	15
28	BTRFS Overview	15
29	BTRFS robustness supporting features	15
30	Cheap, fast, and atomic snapshots and rollback	15
31	Repair and recovery	16
32	Checksumming	16

33 Introduction

34 This design identifies circumstances that, though undesired because of the risk
35 of loss of functionality, cannot be completely avoided and provides suggestions
36 for dealing with them in such a way that as little functionality as possible is
37 lost.

38 Note that improving D-Bus' robustness is a topic that will be covered in a later
39 stage in its own design document. About securing D-Bus services, please see
40 the security design.

41 **Requirements**

42 Minimize loss of data and loss of functionality due to data corruption in these
43 abnormal circumstances:

- 44 • Unexpected power loss
- 45 • Unexpected removal of storage devices
- 46 • Unexpected lack of disk space
- 47 • Physical damage to the media and other hardware errors

48 Minimize loss of functionality due to processes hogging these shared resources:

- 49 • CPU
- 50 • GPU
- 51 • I/O
- 52 • memory
- 53 • network queue
- 54 • D-Bus daemon

55 **Approach**

56 This section explains how to address the requirements in several specific cases,
57 taking into account different data sets and circumstances.

58 **Application data**

59 This section contains recommendations about how to robustly deal with data
60 generated by applications.

61 **General guidelines**

62 No software should assume that opening files will always succeed. Failure condi-
63 tions should be dealt with and the process will either continue running with as
64 little loss of functionality as possible, or will log a message and exit. Programs
65 should do the same when writing data (the filesystem may be full, or any other
66 mode of error might occur).

67 For example, if the browser application finds out at start up that the cookies
68 file is corrupted, it should move the old file away (or just delete it) and run as

69 usual other than past persistent cookies will have been lost. Or if there was
70 an error when writing a new persistent cookie to disk, the browser would keep
71 running with that cookie being transient (in memory only).

72 In order to reduce the effects of data corruption, regardless of the causes, it
73 would make sense to store different data sets in separate files. So that if cor-
74 ruption happens in, for example, the browser cookie store, it would not affect
75 unrelated functionality such as playlists.

76 For big data sets, Collabora recommends SQLite with either Write-Ahead Log-
77 ging (WAL¹) or [roll-back journal](#)². For smaller data sets, a robust method is
78 to write to a temporary file and rename it on top of the old one once finished.
79 This method is called “atomic overwrite-by-rename” and is mostly used when
80 editing a file in-place.

81 POSIX requires the atomicity of [overwrite-by-rename](#)³. Btrfs, Ext3 and Ext4
82 give atomic overwrite-by-rename guarantees, as well as atomic truncate guaran-
83 tees. The FAT filesystem guarantees neither.

84 SQLite

85 For applications using SQLite for their storage, Collabora recommends using
86 either WAL or the rollback journal so that transactions are committed atomi-
87 cally. In addition, filesystem-specific tuning would be done by configuring the
88 SQLite system library for optimal performance.

89 WAL will be the best option in most cases, except when transactions will be
90 very big (involving more than 100 MB) and when writes are very seldom, then
91 the rollback journal would be preferred.

92 Collabora will run the [TCL test harness](#)⁴ for SQLite in LAVA, to detect any
93 issues in the specific configuration and software in the target platform. These in-
94 clude robustness tests that reproduce out-of-memory errors, input/output errors
95 and abnormal termination (crashes or power loss).

96 Tracker

97 Tracker stores data in SQLite files, so the robustness considerations that apply
98 to SQLite apply to Tracker as well. By default it uses WAL instead of the tradi-
99 tional rollback journal, which gives better performance for Tracker’s workload
100 with the same robustness guarantees.

¹<http://www.sqlite.org/draft/wal.html>

²<http://www.sqlite.org/draft/lockingv3.html#rollback>

³<http://pubs.opengroup.org/onlinepubs/009695399/functions/rename.html>

⁴<http://www.sqlite.org/testing.html#tcl>

101 **User settings**

102 For configuration settings in general, Collabora recommends using the [GSet-](#)
103 [tings](#)⁵ API from GLib with the [dconf](#)⁶ backend. When updating the database,
104 dconf will write the whole new contents to a new file, then atomically renaming
105 it on top of the old one.

106 For bigger pieces of data (individual settings whose data component exceeds 1
107 KB), Collabora recommends using plain files via a known-robust file-handling
108 library (such as [GKeyFile](#)⁷ from Glib, which is already a dependency) or SQLite.

109 **Media**

110 For media, the meta-data is stored in Tracker, with the actual data files in the
111 /home filesystem and in attached removable devices.

112 If the Tracker database that contains the meta-data has been corrupted, it
113 should be moved to the side (or deleted) and recreated again by indexing all
114 available media files. To minimize the chances of corruption, refer to [Tracker](#).

115 Software that reads the actual media files should assume media files may contain
116 invalid data and ignore them without further loss of functionality. Corrupted
117 media files should not be displayed in the UI.

118 **Caches**

119 All software that uses a cache file should be ready to find that the cache is
120 unusable and cope with it without loss of functionality (temporary degradation
121 of performance is obviously expected in this case though the mechanism by
122 which the cache became corrupted will be treated by developers as a bug to
123 fix).

124 For example, if during start-up the Folks caches are found to be unreadable, lib-
125 folks would remove the corrupted cache files and recreate them, taking a longer
126 time to reply to queries. As the application using Folks would be executing
127 the queries asynchronously, the UI would keep being functional while the query
128 executes.

129 Examples of other components that use caches and that should cope with cache
130 corruption are the browser and the email client.

131 **Filesystems**

132 The reliability with which data is stored depends on both the storage medium
133 as well as the filesystem. In this section, we cover FAT32 and Btrfs. Ext4 is
134 mentioned, as it is a popular default filesystem on many Linux distributions –

⁵<http://developer.gnome.org/gio/stable/GSettings.html>

⁶<https://wiki.gnome.org/Projects/dconf>

⁷<https://developer.gnome.org/glib/stable/glib-Key-value-file-parser.html>

135 however it doesn't suit the needs of the rollback system – either for system roll-
136 backs (See the System Update and Rollback Design) or for application rollbacks
137 (See the Applications Design).

138 The FAT32 filesystem is not robust under abnormal circumstances since it was
139 not made for devices which could be disconnected at any moment. In general, an
140 approach where writes to the device are tightly controlled and restricted to small
141 time-windows would help minimize the chances of corruption. See the *Media
142 and Indexing* design for a detailed explanation of the issues and suggestions.

143 The Ext4 filesystem is quite robust under power failure by default. It can be
144 made even more robust by [mounting it under data=journal](#)⁸ mode, but at a
145 large cost to performance.

146 Btrfs has been created on very robust principles, building upon the experience
147 of Ext4. Some brief technical details are provided at the end of this document
148 in [BTRFS overview](#).

149 Filesystem options

150 Filesystems usually have parameters that can be tuned to suit specific work-
151 loads. Some of them affect performance as well as robustness; either by trading
152 off between the two, or by taking advantage of specific hardware features avail-
153 able with the storage media.

154 • **FAT32** is a simple filesystem that does not have many filesystem options
155 related to performance or robustness. Since we will not be creating any
156 FAT32 partitions ourselves, only mount-time options are interesting for
157 us. The recommended options are listed below:

158 – sync, flush These filesystem options ensure that the kernel, as well as
159 the filesystem, flush data to the partition as soon as possible. This
160 greatly reduces the chances of data loss or filesystem corruption when
161 USB drives are yanked out by the user.

162 – ro (read only) It is recommended that FAT32 partitions be mounted
163 read-only to avoid filesystem corruption, and other related problems
164 as detailed in the “*Media and Indexing Design*” in the section “*In-
165 dexing database on removable device*”.

166 • **Btrfs** is relatively new, and so does not have many options relevant to
167 our needs of enhancing reliability on eMMC storage media. The available
168 options are listed below.

169 – *Mount-time options:*

170 * commit=number (default: 30) Set the interval of periodic com-
171 mit. This option is recent ([since kernel 3.12](#))⁹.

⁸<http://kernel.org/doc/Documentation/filesystems/ext4.txt>

⁹https://btrfs.wiki.kernel.org/index.php/Mount_options

172 * `ssd` This option enables SSD-specific optimizations and disables
 173 some optimisations specifically for rotating media. This option
 174 is enabled automatically on non-rotating storage.

175 * `Recovery` (default: off) This option can be used to attempt re-
 176 covery of a corrupted filesystem (See [Repair and recovery](#)).

177 – *Filesystem creation options:*

178 * `-s` `sector-size` This is the size of the filesystem blocks used for
 179 allocations. Ideally, this should be the same size as the block
 180 size for the storage medium.

181 * `-M`

182 This sets BTRFS to use “mixed block groups” - a mode that
 183 stores data and metadata chunks together on disk for more effi-
 184 cient space utilization for small filesystems – but incurs a perfor-
 185 mance penalty on large ones. This option is not mature and will
 186 be evaluated in the future.

187 The System Updates and Rollback Design describes the partition layout for
 188 Apertis. Not all the partitions have the same requirements, so both the FAT32
 189 and BTRFS filesystems are used. The partitions are configured as:

- 190 • **Factory Recovery** – This partition is never mounted read-write and must
 191 be readable by the boot loader. Currently the boot loader for Apertis
 192 – U-boot – does not support BTRFS. While patches exist to add that
 193 functionality, they have not yet seen widespread testing. FAT32 will likely
 194 be the filesystem chosen for the factory recovery image.
- 195 • **Minimal Boot partitions** – These partitions must also be readable by
 196 the boot loader, and are currently FAT32. They are not normally mounted
 197 at run-time, instead they are created, mounted, and populated by the
 198 system update software once – and only ever accessed by the boot loader
 199 afterwards. They will be mounted with the “sync” and “flush” flags.
- 200 • **System** - Since BTRFS provides an excellent snapshot mechanism to
 201 assist system rollbacks (See [Cheap, fast, and atomic snapshots and roll-](#)
 202 [back](#)), this partition will be populated with a BTRFS filesystem created
 203 with the appropriate sector size for the storage device. It may be created
 204 with mixed block groups to save storage space if that option does not lead
 205 to instability. It will be mounted with the `ssd` option as well as read-only.
 206 During a system update a single subvolume of the system subvolume will
 207 be mounted read-write. The repair mount option will never be attempted
 208 on the system partition, instead rollbacks or factory recovery will be used
 209 to avoid potentially putting the system into an unknown state.
- 210 • **General Storage** – This partition shares similar requirements to the
 211 system partition. It will be BTRFS, created with an appropriate sector
 212 size and possibly mixed block groups. It will be mounted with the `ssd`

213 option. This is the only built-in non-volatile storage that will always be
214 mounted read-write. In the case of a damaged filesystem, repair may be
215 attempted on this partition.

216 Additionally, there are 2 partitions for raw status flag data that do not use
217 filesystems at all. See the System Updates and Rollback Design for more details.

218 Checksumming

219 Checksumming is used for detecting filesystem corruption due to any reason.
220 Different filesystems have different mechanisms for checksumming which give
221 us coverage for various different causes of filesystem corruption. Each mecha-
222 nism consumes I/O and CPU resources, and that must be weighed against the
223 advantages that it gives us.

224 It is important to note that checksumming does not protect us against corrup-
225 tion or help us in fixing the root cause of the corruption; it only allows us
226 to detect filesystem corruption when it happens. Hence, it is only useful as a
227 warning sign and recovering from data corruption is beyond the scope of this
228 feature.

- 229 • **FAT32** is a very old and simplistic filesystem, and it has no inbuilt facili-
230 ties for checksumming.
- 231 • **Btrfs** maintains a *checksum tree* for all the blocks that it allocates and
232 writes to. Hence, all file data and metadata is checksummed. This is the
233 default behaviour and the current checksum algorithm uses few resources.
234 This method of checksumming can detect all the ways in which corruption
235 can occur to data on the filesystem. See [Checksumming](#) for more detail.
- 236 • **Ext4** maintains checksums for journal data only, no checksumming of file
237 data takes place.

238 Alignment

239 The first piece of tuning that a filesystem on flash storage needs, is a proper
240 mapping of the filesystem blocks to the page size of the erase blocks on the flash.
241 This consists of two parts:

- 242 1. Ensuring that the filesystem and storage erase block sizes match using
243 filesystem creation options.
- 244 2. Aligning the block allocations in the filesystem with the storage blocks
245 by using the appropriate offsets while partitioning, or while creating the
246 filesystem.

247 If either of these is not satisfied, each filesystem block write will trigger two or
248 more flash block writes, and reduce the performance as well as reliability of the
249 MMC card.

250 The storage erase block size [can be read](#)¹⁰ from /proc/mtd or from U-Boot but
251 the flash storage can report something different than the real numbers. Linaro-
252 image-tools is now able to generate images with a correct alignment.

253 **Testing**

254 Collabora will add tests to LAVA for testing how FAT32 and Btrfs behave on
255 the i.MX6 under stress, as well as for tuning the above mentioned parameters
256 for reliability and performance.

257 **Root filesystem**

258 The approach will be to mount as many parts of the root filesystem read-only
259 as possible such that the only writes to it would be during updates. This would
260 reduce the chances of catastrophic filesystem corruption in the event of power
261 failure and invalid system file modification by bugs in system or application
262 software. The only partition that is to be mounted writable is the user partition
263 that will be mounted in /home. All the other writable parts of the / filesystem
264 will be backed by tmpfs, located in RAM. We will avoid the lack of space
265 problem by only storing small files in tmpfs or files which don't take space (lock
266 files, socket files). Bigger files such as programs, libraries, configuration files will
267 remain on disk and available read-only.

268 See the *System Updates and Rollback* design for detailed information about the
269 robustness of the update process.

270 **Other filesystems**

271 The system should be able to function even if mounting one or more of the
272 non-essential file systems fails. Even if the system is able to keep running, it
273 would do so with reduced functionality, so some recovery action would need
274 to be taken in order to regain the lost functionality. The system should try
275 to recover automatically as far as possible. In the case of unrecoverable system
276 failure, the user can be instructed at system boot to request technical assistance
277 at a service shop.

278 **Main storage**

279 In case of power loss, the flash media can become corrupted due to how writes
280 are performed. Apertis will be notified via a GPIO signal 100 milliseconds before
281 power is completely lost, in order to give the flash controller time to commit to
282 non-volatile media what is in its cache.

283 Given the short time available and the general slowness of flash devices when
284 writing, we recommend that the signal is handled in the kernel, because
285 userspace will not have enough time to react (depending on the load and the

¹⁰<https://bootlin.com/blog/managing-flash-storage-with-linux/>

286 scheduler, it could take from 10 ms to 100 ms for the signal to start being
287 processed by a userspace process). A device driver should be written that,
288 when the GPIO signal is received:

- 289 1. stops flushing dirty pages to the drive,
- 290 2. tells the flash controller to flush its caches to permanent storage, and
- 291 3. starts the shutdown sequence.

292 The device driver will start handling the signal 10-100 μ s after the GPIO is
293 activated. In spite of this, if the device has big caches and is slow to write,
294 corruption of arbitrary data blocks can still happen.

295 In general, drive health data should be monitored so that the user can be notified
296 about disk failures which require a garage visit for hardware replacement.

297 As no more dirty pages will be flushed to the storage device when the GPIO
298 signal is received, the data in the page cache will be lost. To reduce the amount
299 of data that could be lost, eMMC reliable writes can be used, and the page cache
300 configuration can be tuned. But it has to be noted that use of reliable writes
301 and reducing the amount of in-flight data is a trade-off against performance
302 that can be quantified only on the final hardware configuration through direct
303 experimentation.

304 **Removable devices**

305 External devices that can be removed at any moment are not reliable for writ-
306 ing of critical data. In addition to the problem of corruption of files being
307 written, wear leveling by the controller might corrupt unrelated blocks which
308 might even contain the directory table or the file allocation table, rendering the
309 whole partition unusable.

310 The quality of external storage devices such as flash drives varies greatly, in some
311 cases the device will unexpectedly stop responding to commands, or data will
312 be lost. Applications that write to removable drives must be robust enough to
313 be able to continue in the face of such errors with minimal loss of functionality.

314 As mentioned in [Filesystems](#), the safest way to use removable drives is by re-
315 stricting the processes that can write to the drive, and minimizing the time-
316 window for the writes. For that to be practical, there should be a system ser-
317 vice that is the only one allowed to write to removable devices and that would
318 accept requests from applications, remount the device read-write, write the new
319 contents, then remount read-only again.

320 Since, for interoperability reasons, the filesystem used in removable devices is
321 FAT32, in addition to the issues mentioned in this section, the robustness con-
322 siderations that were explained earlier in [Filesystems](#) also apply.

323 Mitigating the effects of lack of disk space

324 In order to reduce the chances that the system will find itself in a situation
325 where lack of disk space is problematic, it is recommended that available disk
326 space is monitored and applications notified so they can react and modify their
327 behavior accordingly. Applications may chose to delete unused files, delete or
328 reduce cache files or purge old data from their databases.

329 The recommended mechanism for monitoring available disk space is for a dae-
330 mon running in the user session to call *statvfs* (2) periodically on each mount
331 point and notify applications with a D-Bus signal. [Example code](#)¹¹ can be
332 found in the GNOME project, which uses a similar approach (polling every 60
333 seconds).

334 Additionally, so error messages can be stored also in low-space conditions, it
335 is recommended that *journald* is configured to leave an amount of free space
336 smaller than the reserved blocks of the filesystem that backs the log files. This
337 way, applications will still be able to log messages after applications have con-
338 sumed all the sace available to them.

339 In case applications cannot be trusted to properly delete non-essential files, a
340 possibility would be for them to state in their manifest where such files will be
341 stored, so the system can delete them when needed.

342 In order to make sure that malfunctioning applications cannot cause disruption
343 by filling filesystems, it would be required that each application writes to a
344 separate filesystem.

345 It may be worth noting that temporary directories should be emptied on reboot.

346 Resource management

347 The robustness goal of resource management is to prevent one or more applica-
348 tions from disrupting basic functionality due to excessive resource consumption.
349 The basic mechanism for this is to allocate resources in such a way that applica-
350 tions cannot starve services in the base system. This is to be achieved firstly by
351 changing the resource allocation policy to give higher priority to services, and
352 secondly by limiting the maximum amount of resources that an application can
353 consume at a time.

354 Resource limits are capable of helping ensure a process does not render the
355 whole system unresponsive. However, some design decisions also play an im-
356 portant role here. If the user has no way to kill the process that became too
357 slow or unresponsive, the user experience will suffer. The same goes for the
358 case in which an application gets stuck into a failing scenario, such as a web
359 browser automatically loading pages that were open when the browser closed
360 unexpectedly. For these reasons care must be exercised while designing the user

¹¹<http://git.gnome.org/browse/gnome-settings-daemon/tree/plugins/housekeeping/gsd-disk-space.c#n693>

361 interactions for both the system chrome and applications to be sure such cases
362 are addressed.

363 If, despite throttling, some processes still impact the overall user experience
364 negatively because of excessive resource usage, there is the option of identifying
365 those processes and terminating them. Collabora recommends against this be-
366 cause it is very difficult to automatically distinguish between processes that use
367 large amounts of resources due to malfunction or maliciousness and processes
368 that use excessive resources for legitimate purposes. Killing the wrong process
369 may free up resources but is likely to be perceived by the user as a severe defect
370 in the overall user experience.

371 As a general recommendation, for optimal responsiveness, applications should
372 not block the UI thread when calling anything that is not assured to return
373 almost immediately, which includes all local or remote I/O operations. When
374 the potential duration of an operation is a considerable portion of the commonly-
375 considered maximum acceptable response time (100 ms), it should be done
376 asynchronously. GLib contains [asynchronous APIs](#)¹² for I/O in its [file](#)¹³ and
377 [streaming](#)¹⁴ classes.

378 CPU

379 To make sure that important processes have available CPU cycles even when mal-
380 functioning or malicious applications monopolise the CPU, it is recommended to
381 set task scheduler priorities according to the importance of processes. Systemd
382 can do this for services by setting the [CPUSchedulingPriority](#)¹⁵ property in the
383 service unit file of the process. When the process described by the service unit
384 file starts new processes, they stay in the same cgroups and they keep the same
385 CPUSchedulingPriority.

386 At present (Q1 2014), systemd manages the user session on target images but
387 not on the SDK. With the user session managed by systemd, the priorities of ap-
388 plications are no longer set by the application launcher using [sched_setscheduler](#)
389 [\(2\)](#)¹⁶.

390 If there are processes that need real-time capabilities, or that should have very
391 low CPU access, the CPUSchedulingPolicy property can be used to change to
392 the rr (real-time) or idle scheduling policies. Real-time access for a process
393 should be carefully considered and tested because it can have a negative impact
394 on the process and even the entire system.

395 For identifying processes that use an excessive amount of CPU, the [cpuacct](#)¹⁷
396 cgroups controller can be used.

¹²<http://developer.gnome.org/gio/stable/async.html>

¹³http://developer.gnome.org/gio/stable/file_ops.html

¹⁴<http://developer.gnome.org/gio/stable/streaming.html>

¹⁵<http://0pointer.de/public/systemd-man/systemd.exec.html>

¹⁶http://www.kernel.org/doc/man-pages/online/pages/man2/sched_setscheduler.2.html

¹⁷<https://www.kernel.org/doc/Documentation/cgroup-v1/cpuacct.txt>

397 Though it is not recommended to automatically terminate local applications
398 with excessive CPU usage, it makes sense for web pages. Web pages are not
399 screened before they execute on the system, hence it is important to ensure that
400 their ability to disrupt system functionality is minimised. For this, WebKit can
401 detect when a block of JavaScript code has been executing for too long, pause
402 it, and give the embedding application the possibility of canceling the execution
403 of this block of code. Collabora has added API to WebKit-Clutter for this.

404 I/O

405 Similar to CPU usage, Collabora recommends giving priority to important pro-
406 cesses when there is contention for I/O bandwidth. Collabora recommends that
407 important services have a value for the property `IOSchedulingPriority` lower
408 than 4 (the default). If, for any reason, some applications need priorities other
409 than the default, the application launcher can use the `ioprio_set`¹⁸ (2) syscall
410 to change their priority. When the process described by the service unit file
411 starts new processes, they stay in the same cgroups and they keep the same
412 `IOSchedulingPriority`.

413 Memory

414 Collabora recommends putting a single limit on the amount of memory that the
415 whole application set can allocate so a fair reserve is left for the base software.
416 This limit should be just big enough so that the Apertis instance never reaches
417 the “out of memory” (OOM¹⁹) condition at the system level. For example, if
418 the total of memory available for processes is 1GB, there is no swap, and we
419 know that the services in the base system should need a maximum of 300MB,
420 then all applications should belong to a cgroup that is limited to 700MB of
421 memory.

422 In specific cases, it may make sense to put a different limit on a specific appli-
423 cation, but it can easily be counterproductive and cause a waste of memory.

424 Something else worth doing is to make sure that the OOM killer²⁰ selects ap-
425 plications for killing instead of system services. For this, the systemd prop-
426 erty `OOMScoreAdjust` can be used to reduce the chances that a service will be
427 killed. For applications, it is recommended that the application launcher sets
428 its `/proc/<pid>/oom_score_adj` (see [here](#)²¹) to be higher than 0. The ideal
429 value may vary depending upon the importance of each application.

430 With the example setup mentioned before, the OOM killer will terminate the
431 bulkiest application when one of these conditions are met:

¹⁸http://www.kernel.org/doc/man-pages/online/pages/man2/ioprio_set.2.html

¹⁹http://en.wikipedia.org/wiki/Out_of_memory

²⁰<http://lwn.net/Articles/317814/>

²¹<http://www.kernel.org/doc/Documentation/filesystems/proc.txt>

- 432 • The total memory taken by applications all together is going to increase
433 over 700MB.
- 434 • The total memory taken by all processes (services plus applications) is
435 going to increase over 1GB.

436 To make better use of the available memory, it's recommended that applications
437 listen to the cgroup notification `memory.usage_in_bytes`²² and when it gets
438 close to the limit for applications, start reducing the size of any caches they
439 hold in main memory. It may be good to do this inside the SDK and provide
440 applications with a glib signal that they can listen for.

441 Network queue

442 Processes would be classified into cgroup classes such as:

- 443 • Interactive (VoIP, internet radio)
- 444 • Semi-interactive (web pages, maps)
- 445 • Asynchronous (mail, app notifications, etc)
- 446 • Bulk (downloads, system updates)

447 Cgroup controllers are only used for classification of outgoing packets. `NET-`
448 `PRIO_CGROUP`²³ and `NET_CLS_CGROUP`²⁴ would be used for setting the
449 priority, and for classifying processes into cgroups. By thus tagging packets
450 with the cgroup of applications and services, `tc`²⁵ can be used to set limits to
451 the rate at which processes send packets (<http://lartc.org/howto/>).

452 Bandwidth rate-limiting would be required to ensure interactive streams do not
453 get starved by lower priority streams.

454 There is little we can do about latency for applications like VoIP, since even when
455 the bandwidth is sufficient, the bottlenecks are the hardware buffers, queues,
456 and scheduling on various devices outside the control of our system. This is an
457 open problem in networking, and a large part of it is related to `Bufferbloat`²⁶.

458 Note that there's no robustness issue that can be prevented by limiting the rate
459 at which processes receive incoming packets.

460 GPU

461 As explained in the WebGL design, the `GL_EXT_robustness`²⁷ extension pro-
462 vides a mechanism by which the watchdog in the GL implementation can reset

²²<http://www.kernel.org/doc/Documentation/cgroup-v1/memory.txt>

²³<http://lwn.net/Articles/474695/>

²⁴http://docs.fedoraproject.org/en-US/Fedora/16/html/Resource_Management_Guide/sec-net_cls.html

²⁵<http://lartc.org/manpages/tc.txt>

²⁶<http://en.wikipedia.org/wiki/Bufferbloat>

²⁷http://www.khronos.org/registry/gles/extensions/EXT/EXT_robustness.txt

463 the GPU, invalidating all GL contexts and thus stopping all GPU activity.

464 Unfortunately, this only prevents denial of service (DoS) conditions caused by
465 WebGL, because processes must opt-in to use this extension. Thus, applications
466 may intentionally or unintentionally ignore the extension and continue monopolizing the GPU. Within the web browser, scripts that use WebGL and take over
467 the GPU will be interrupted and terminated by the browser.
468

469 If it runs its own GL implementation, then it could monitor GPU resource
470 usage and reset those contexts that seem to be disrupting the rest of the system. It could notify processes via the `GL_EXT_robustness` extension and even
471 terminate them if they ignore the context reset notifications.
472

473 **Accounting**

474 Besides setting limits on resources, cgroups also allows to retrieve resource usage metrics. As examples, for CPU usage the *cpuacct* cgroup controller contains the *usage*, *stat* and *usage_percpu* reports; the *memory* controller provides usage data in its *stat* report; the *blkio* controller has *throttle.io_serviced* and
478 *throttle.io_service_bytes*.

479 **USB undervoltage**

480 In the case that the system momentarily isn't able to power connected USB devices such as MP3 players or smartphones due to voltage drops, the system will
481 power off and on again these devices, so that the connection gets reestablished
482 and the user experience gets affected as little as possible.
483

484 **Risks**

- 485 • FAT32 is fundamentally unreliable, specially on removable devices.
- 486 • Robustness of flash media varies greatly and the user may not be able
487 to distinguish failures caused by the hardware from failures due to the
488 software.
- 489 • Excessively-low resource limits for applications can lead to resource waste;
490 excessively-high may be less effective in avoiding DoS. There may not exist
491 a good middle point.
- 492 • Heuristics used to determine when to kill a process with excessive resource
493 usage are not perfect and can cause major failure from the user point of
494 view.
- 495 • If Vivante does not implement `GL_EXT_robustness` properly, web pages
496 could DoS the whole system.

- 497 • Bugs in the OpenGL implementation can lead to instability, data loss and
498 privacy breaches that can be triggered from web pages.
- 499 • If the flash media loses power while a block is open for writing, it is possible
500 that several random blocks elsewhere in the same drive will be corrupted.
501 This can affect other filesystems, even if they are mounted read-only.

502 Design notes

503 The following items have been identified for future investigation and design work
504 later in the project and are thus not addressed in this design:

- 505 • Vulnerability to DoS attacks in D-Bus and proposed solutions.
- 506 • Optimization of the SQLite configuration parameters for the specific
507 filesystems in use in Apertis.

508 No updates as of March 2014.

509 BTRFS Overview

510 The most powerful feature of [Btrfs](#)²⁸ is the fact that all information (data +
511 metadata) is stored in the same basic data structures, and all modification of
512 these data structures is performed in a copy-on-write (CoW) fashion.

513 Since all information on disk is stored using the same type of data structure, this
514 allows metadata and data to share features such as checksumming and striping.

515 This combined with the fact that Btrfs uses CoW while modifying all informa-
516 tion, means that in theory, the filesystem is always consistent if the storage
517 device supports “[Force Unit Access](#)”²⁹ correctly. However, in practice, fileys-
518 tem bugs, a lack of maturity in the code, and other (unforeseen) problems may
519 prevent this.

520 BTRFS robustness supporting features

521 Cheap, fast, and atomic snapshots and rollback

522 All snapshots in Btrfs are CoW copies of the subvolume being snapshotted with
523 an incremented reference count for the blocks. As a result, creating snapshots
524 is very fast, and they take up a negligible amount of space. Just like every
525 other operation, the snapshot is created atomically by the use of transactions
526 and sequenced flushes. Further, all snapshots are actually just subvolumes, and
527 hence can be mounted on their own.

²⁸<https://btrfs.wiki.kernel.org/>

²⁹[https://en.wikipedia.org/wiki/Disk_buffer#Force_Unit_Access_\(FUA\)](https://en.wikipedia.org/wiki/Disk_buffer#Force_Unit_Access_(FUA))

528 Unlike LVM2, which creates snapshots in the form of block devices that can be
529 mounted, Btrfs creates snapshots in the form of subvolumes, which are repre-
530 sented as subdirectories.

531 Even though snapshots are displayed in a subdirectory they are not “owned” by
532 that subvolume. Snapshots and subvolumes are identical in Btrfs, and are first-
533 class citizens with respect to other subvolumes. This means that the default
534 subvolume can be set at any time. The change will be made the next time the
535 filesystem is mounted. All the subvolumes, except the top-level subvolume, can
536 also be deleted; irrespective of their relationships with each other.

537 **Repair and recovery**

538 If, for any reason, the root node or the superblock gets corrupted and the filesys-
539 tem cannot be mounted, mounting in recovery mode will make btrfs check the
540 superblock (or alternate superblocks if the superblock is also corrupted) for
541 alternate roots from previous transactions. This is possible because all modifi-
542 cations to the Btrfs trees are done in a CoW manner and existing roots are not
543 deleted. The filesystem stores the last four roots as a backup for the recovery
544 option.

545 **Checksumming**

546 The header of every chunk of space in Btrfs has space for 32-bytes of check-
547 sums of the chunk itself. In addition, there is a checksum tree which maintains
548 checksums for each block of data. Since the data as well as the metadata blocks
549 are referenced in the checksum tree, all information in the filesystem is check-
550 summed.

551 Currently, Btrfs uses the CRC-32 checksum algorithm, but there are plans to
552 upgrade that, and add the option to set the checksum algorithm when the
553 filesystem is created.