



Compositor security

1 Contents

2	Use-cases	3
3	Home screen	3
4	Platform UI elements	3
5	Trusted output	4
6	Launching a program	4
7	Last-used mode	6
8	Main window selection	7
9	Child windows	8
10	Notifications	9
11	Focus-stealing	10
12	Non-graphical programs	11
13	Screenshots	11
14	Synthesized input	12
15	Trusted input paths	12

16 Further reading 13

17 The *compositor* is the component of Apertis that is responsible for drawing
18 application windows and other graphical elements on the screen.

19 In Apertis the compositor runs as the [agl-compositor](#)¹ executable from a **sys-**
20 **temd user unit** and launches **maynard** as desktop manager. This is a thin
21 executable wrapper around the library **libweston**, which provides the majority
22 of its functionality.

23 In Wayland, the compositor *is* the display server. Graphical programs arrange
24 for their graphics to be displayed by creating a buffer (a *surface*) in GPU mem-
25 ory, drawing their text, images etc. into that buffer, then sending requests to
26 the Wayland compositor which ask the compositor to include that surface in
27 the final 2D scene. Unprivileged programs cannot display graphics until the
28 compositor is ready, so we can be sure that the compositor's policies are applied
29 to every surface.

30 We aim to provide the usual security properties described in the [Security design](#)
31 [document](#)²:

- 32 • confidentiality
- 33 • integrity
- 34 • availability

35 for the two mechanisms provided by the compositor:

- 36 • output (placing application windows on the screen)

¹https://docs.automotivelinux.org/en/master/#5_Component_Documentation/1_agl-compositor/

²https://www.apertis.org/concepts/archive/application_security/security/

- input (dispatching input events such as touchscreen touches and gestures to applications)
- [Wayland Compositors - Why and How to Handle Privileged Clients] provides a good overview of how those security properties apply to compositors.

Use-cases

“The platform” refers to the overall Apertis platform, including the compositor, application manager and so on.

Because we anticipate that the desired graphical presentation and user experience (UX) will be a point of differentiation for OEMs, each of these requirements should be interpreted as a requirement that it is possible for the platform to behave as specified, and a recommendation that OEMs’ platform variants should do so unless it conflicts with their desired UX. For example, for brevity, we will use “the compositor must …” as shorthand for “it must be possible for the compositor to …”, and we recommend that OEMs’ compositors should have that behaviour unless it conflicts with their desired UX”.

Home screen

In some circumstances, such as when the Apertis device is switched on for the first time, it must go into a default state.

- The platform must draw a “home screen” or launcher from which further programs can be launched.
- The home screen may either be part of the compositor, or a separate graphical program.
- Pressing a button or menu entry representing an application [entry point](#)³ results in the relevant graphical program being started.

(These are aspects of input and output availability).

On Apertis, this job is accomplished by **maynard** which is the default desktop manager.

Platform UI elements

In addition to the home screen, there might be UI elements which are outside the scope of any particular application window, such as a status bar, **notifications**, **system-modal dialogs**, or the UI controls used for **application-switching**.

- The OEM-specific visual design might reserve regions of the screen for these visual elements. We recommend that this is done.

³https://www.apertis.org/concepts/archive/application_framework/application-entry-points/

- 70 – For example, the equivalent features in Android are the small re-
71 region at the top of the screen that is normally reserved for the status
72 bar, and the larger region at the bottom or side of the screen that
73 is normally reserved for the navigation bar (Back, Home and Apps
74 buttons).
- 75 • The compositor may either draw each of those UI elements itself, or ar-
76 range for separate programs to provide them.
- 77 • Some of these UI elements must remain visible at all times (they must be
78 displayed on top of ordinary program windows), unless the compositor’s
79 UX calls for them to be hidden under certain specific circumstances.
 - 80 – For example, Android allows applications to request that the status
81 bar and navigation bar are hidden, but the gestures to reinstate them
82 are always available, and the operating system displays a reminder
83 of those gestures when they become hidden.
- 84 • If separate programs provide some or all of these UI elements, then normal
85 platform startup must arrange for them to be launched.

86 *(These are aspects of input and output availability).*

87 The approach used in **agl-compositor** is to support this by providing protocol
88 extensions. Thanks to them, surfaces can have roles, such as popup, fullscreen,
89 split_vertical or split_horizontal and it is also possible to configure them as
90 panels to be always visible and anchored to one of the edges. On Apertis,
91 **maynard** implements these protocol extensions to display the UI elements.

92 **Trusted output**

- 93 • The compositor must not allow unprivileged programs to display their
94 content in the regions of the screen that are reserved for these UI elements,
95 unless the compositor’s UX design specifically allows it. This is a *trusted*
96 *path* with which the platform can display information to the user. (*Output*
97 *integrity*)
 - 98 – Ideally, the APIs provided to programs should be designed so that it
99 is impossible to request display in a forbidden area.
 - 100 – If the APIs provided to programs are such that the program can
101 attempt to display in these regions, and an unprivileged program
102 attempts to do so, this must be detected and prevented.

103 Since **agl-compositor** is a Wayland compositor, applications cannot request a
104 specific region to display their content. It is the responsibility of the compositor
105 to choose the appropriate place while enforcing its policies.

106 **Launching a program**

107 When a graphical program is launched, after carrying some non-graphical ini-
108 tialization, it will create a surface, fill it with the first frame that it wants to be
109 displayed, and submit that surface to the compositor for display.

- 110 • The compositor must be able to identify that surface as having come from
111 that graphical program. In particular, it must be able to determine the
112 [app-bundle](#)⁴ and [user account](#)⁵ that originated the surface. (*Input and*
113 *output integrity*)
 - 114 – *Non-requirement*: If an app-bundle is allowed to contain multiple
115 graphical programs, the ability to distinguish between those graphical
116 programs is optional. We treat the app-bundle as a security boundary,
117 but we do not place a security boundary between individual graphical
118 programs within an app-bundle.
- 119 • This identification must be securely authenticated. If a different user
120 account or app-bundle asks to display a surface, one of these options must
121 be true:
 - 122 1. (Preferred) The compositor obtains the originating program’s user
123 account and app-bundle directly from the Linux kernel or some other
124 trusted platform component, and there is no opportunity for the
125 originating program to give false information.
 - 126 2. The originating program tells the compositor which user account and
127 app-bundle it claims to be, and the compositor verifies in a secure
128 way that this claim is true.
 - 129 – *Non-requirement*: If an app-bundle is allowed to contain multiple
130 graphical programs and the compositor distinguishes between them,
131 it is acceptable for it to be possible for a graphical program to be able
132 to impersonate a different graphical program in the same bundle.
- 133 • The compositor must perform whatever appropriate smooth graphical
134 transition is desired (for example a cross-fade, animated movement, or
135 a simple atomic change between one frame and the next) between the
136 home screen and the graphical program’s surface as the main contents of
137 the screen.
- 138 • If the compositor’s UX involves multiple tiled content areas, the graphical
139 program must be displayed in the desired content area.
- 140 • If the compositor’s UX involves floating or cascading windows (as seen
141 in GNOME, Windows, etc.), the graphical program must be displayed in
142 the location chosen by the compositor. It may influence that location by
143 setting “hints” in its requests, but the compositor must be free to ignore
144 those hints.
- 145 • The compositor must arrange for any [platform UI elements](#) that should
146 remain visible at all times to remain visible and interactive during this
147 process (*input and output availability*):
 - 148 – if they are provided by the compositor itself, they must be layered
149 above the graphical program’s surfaces in the compositor’s scene-
150 graph;
 - 151 – if they are provided by a separate “shell” program, the surfaces repre-
152 senting them must be layered above the surfaces from the graphical

⁴<https://www.apertis.org/glossary/#app-bundle>

⁵<https://www.apertis.org/glossary/#user-account>

- 153 program.
- 154 • The compositor must deliver location-specific input events such as touch-
155 screen touches to the application at the relevant location, and to no other
156 application. (*Input availability, input confidentiality*)
 - 157 • In particular, if application windows can overlap (for example stacking or
158 cascading), and application A is in front of application B, then application
159 A must not be able to trick the user into entering confidential input that
160 was intended for application B by making itself transparent or almost-
161 transparent, so that the user interface of application B shows through
162 (*clickjacking*⁶). (*Input confidentiality*)
 - 163 • The compositor must deliver non-location-specific input events such as
164 touchscreen edge-swipe gestures to the current application, using a defini-
165 tion of “current” that is part of its UX, and to no other application. (*Input*
166 *availability, input confidentiality*)

167 Thanks to the fact that **agl-compositor** is based on Wayland, an application
168 only controls the contents of its surfaces and the compositor chooses where
169 applications are displayed. That makes input and output availability, as well as
170 input confidentiality easy to ensure.

171 The Wayland protocol operates via an `AF_UNIX socket`⁷, just like D-Bus, so
172 applications can be identified by their AppArmor profile and uid using the same
173 credentials-passing mechanisms that are already available in D-Bus.

174 Also, since user applications are meant to be deployed and launched using the
175 `Apertis application framework`⁸, applying the principles described in `security`
176 `boundaries and thread model`⁹ minimises the risks.

177 Last-used mode

178 In some circumstances, such as when the Apertis device is switched off with
179 a particular app active, UX designers may wish to return to a previous saved
180 state, for example one that was saved during device shutdown (“last-used mode”
181).

- 182 • The platform must arrange for each of the graphical programs that was
183 previously active and visible (in the foreground) to be restarted.
- 184 • When one of those graphical programs asks the compositor to display a
185 surface, the compositor must place it in the same location where it was
186 previously visible.
- 187 • The platform may launch other graphical programs that were running but
188 not visible when the state was saved. They must not become visible until

⁶<https://en.wikipedia.org/wiki/Clickjacking>

⁷<http://wayland.freedesktop.org/docs/html/ch04.html>

⁸https://www.apertis.org/concepts/archive/application_framework/application-framework/#the-next-generation-apertis-application-framework

⁹https://www.apertis.org/concepts/archive/application_security/security/#security-boundaries-and-threat-model

189 the user makes a request to switch to them. Alternatively, the platform
190 may delay starting those graphical programs until the user makes a request
191 to switch to them.

192 (*Input and output availability*)

193 Main window selection

194 The user should have the opportunity to switch between the main (top-level)
195 windows presented by various programs.

196 A graphical program might make it difficult for the user to leave, either acciden-
197 tally (because the program has become unresponsive) or deliberately as a denial
198 of service (because the program is maliciously written or has been compromised
199 by an attacker).

- 200 • The compositor must have the opportunity to intercept input events
201 (touchscreen touches, touchscreen gestures, hardware button presses)
202 regardless of the actions of the program. (*Input availability*)
- 203 • The compositor should always provide a way to return to a home screen
204 or application switcher, from which an unresponsive program can be ter-
205 minated. (*Input and output availability*)
- 206 • The way to return to a home screen or application switcher should be
207 consistent and predictable. For example, Android reserves a small area of
208 the screen for Back, Home and Applications buttons. In older Android
209 versions, applications such as the camera may request that these buttons
210 are displayed unobtrusively, but are not able to hide them altogether; in
211 newer versions, these buttons can be hidden, but the swipe gesture to make
212 them available cannot be disabled, and the user is given a reminder of that
213 gesture which cannot be hidden by the application. (*Input availability,*
214 *output integrity*)
 - 215 – Optionally, specially privileged app-bundles might be given the op-
216 portunity to hide these UI elements, or arrange for one of the app-
217 bundle’s surfaces to be displayed as an overlay “above”them. However,
218 this should be a “red flag”in app-store review, to be granted only to
219 trusted applications.
 - 220 * For example, Android requires the `SYSTEM_ALERT_WINDOW`
221 `permission`¹⁰ for applications that use overlays, and additionally
222 requires that the user has been specifically prompted by the
223 platform to grant this permission to this app.
- 224 • If the compositor receives an input event that it interprets as a request to
225 switch away from the graphical program, for example pressing a “home”or
226 “application switcher”button, then this switch must occur within a reason-
227 able time, even if the current graphical program does not cooperate with

¹⁰<https://developer.android.com/reference/android/provider/Settings.html#canDrawOverlays%28android.content.Context%29>

- 228 that operation. This must have a smooth graphical transition (cross-fade
 229 or animation) if that is the desired UX. (*Input and output availability*)
- 230 – For example, if a bug in the current graphical program results in
 231 it ceasing to respond to messages from the compositor (for example
 232 a deadlock or live-lock situation) and the window switching opera-
 233 tion involves communicating with it, the compositor must not wait
 234 indefinitely for a response. If it gets a response, it may switch imme-
 235 diately; if it does not, it may wait a short time, but after that time
 236 it must continue switching anyway. The maximum wait time should
 237 be chosen so that switching still appears responsive.
 - 238 – Similarly, if the current graphical program is deliberately/maliciously
 239 written with the intention of delaying task-switching as much as pos-
 240 sible, the compositor must still switch within a reasonable time.
 - 241 • Each window offered for switching must be associated with the relevant
 242 app-bundle, for example with a title and/or icon, so that when the user
 243 believes they are switching to a particular window, they can know that
 244 they are in fact switching to a window from the correct trust domain.
 245 (*Input and output integrity*)
 - 246 – The ability to distinguish between windows from different graphical
 247 programs in the same app-bundle is optional, because graphical pro-
 248 grams in an app-bundle share a trust domain.
 - 249 • A UX designer might require a limit on the number of simultaneous win-
 250 dows per app-bundle. For example, an app-bundle might be limited to
 251 having up to 5 entry points in the same or different processes, each with
 252 up to 2 main windows open at any given time.

253 On Apertis **maynard** displays a panel that cannot be hidden which allows to
 254 show the list of available applications and to switch to any of them. It supports
 255 `.desktop` files as source for applications metadata.

256 The compositor enforce timeouts when interacting with surfaces in order to
 257 prevent not responding application to compromise user experience.

258 Child windows

259 A graphical program might include `dialogs`¹¹ in its UX.

- 260 • We recommend that dialogs should normally appear as a direct result of
 261 user activity, but they may also appear as a result of an external event.
- 262 • If the graphical program's corresponding main window is currently dis-
 263 played in a particular location, the dialog should overlay that location.
 264 If the API to open dialogs makes it possible to attempt to place dialogs
 265 elsewhere, and the program does so, the compositor must prevent this.
 266 (*Output integrity*)
- 267 • If surfaces (windows) are tiled, stacked or floating, the dialog may extend
 268 outside the boundaries of the graphical program's main window if desired,

¹¹https://en.wikipedia.org/wiki/Dialog_box

269 but we recommend that this pattern is discouraged. If this is done, it
270 should always be made obvious which surface the dialog belongs to. (*Out-*
271 *put integrity*)

- 272 • The dialog must not prevent the user from [leaving the program], even if
273 it extends outside the main window; in other words, it may be app-modal
274 or document-modal, but must not be system-modal. (*Input and output*
275 *availability*)
- 276 • We suggest encouraging the use of **document-modal dialogs**¹² similar to
277 those in **OS X**¹³ and **GNOME**¹⁴

278 A graphical program might include pop-up or drop-down menus in its UX.

- 279 • Menus typically behave like a document-modal window immediately above
280 their “parent” window.
- 281 • The requirements are essentially the same as for dialogs, although the
282 visual presentation is likely to be different.

283 To enforce these rules, **agl-compositor** only supports top-level surfaces and
284 only allows one main surface per application. Under these restrictions, applica-
285 tions requiring additional surfaces need to create them as child objects of the
286 main one.

287 The support of system popups is implemented through custom protocol exten-
288 sions which can be only used by privileged applications which implement them,
289 like **maynard** does.

290 Notifications

291 External events might result in a *notification*, typically implemented as a “pop-
292 up” window.

- 293 • A calendar might trigger notifications as time passes, for example when
294 an appointment will occur soon.
- 295 • A messaging application (for example email or Twitter) might trigger a
296 notification when new messages are available.

297 These notifications should be displayed by the platform user interface (HMI),
298 either as part of the compositor (like in GNOME Shell) or a separate process.

- 299 • If there is a current notification, the platform should draw a visual rep-
300 resentation of it, displaying it “above” any current window. (*Output avail-*
301 *ability for the notification*)
- 302 • If there is no current notification, any program (including non-graphical
303 programs) may trigger a new notification. (*Output availability for the*
304 *notification*)

¹²https://en.wikipedia.org/wiki/Dialog_box#Document%20modal

¹³https://developer.apple.com/library/mac/documentation/UserExperience/Conceptual/OSXHIGuidelines/WindowDialogs.html#//apple_ref/doc/uid/20000957-CH43-SW2

¹⁴<https://wiki.gnome.org/Design/OS/ModalDialogs>

- 305 • Each notification should be visually associated with the appropriate app-
306 bundle, perhaps via an icon and title. (*Output integrity*)
- 307 • Notifications should be drawn in such a way that only the compositor (or
308 the trusted notification service, if separate) can produce the same visual
309 result, for example by displaying it over the top of **platform UI elements** in
310 a way that would not be possible or would not be allowed for an ordinary
311 application window. (*Output integrity*)
- 312 • There should be a straightforward mechanism by which the driver can close
313 any notification, minimizing distraction. (*Input and output availability for
314 other UI components*)
- 315 • High-priority platform components such as navigation must be able to
316 force their notifications to be displayed instead of, or “above”, other com-
317 ponents’ notifications. (*Output availability for the higher-priority notifica-
318 tion*)
- 319 • Excessive notifications by an application might be distracting. The com-
320 positor must have the opportunity to limit the number of notifications
321 per app-bundle or deny notification display altogether, with an optional
322 user-configurable limit per application so that the user could selectively
323 silence an app-bundle that they found distracting.
- 324 • The precise handling of notifications (for example topics such as how mul-
325 tiple simultaneous notifications are handled) is outside the scope of this
326 document.
- 327 • If the notification has “actions”, for example a button to go to the relevant
328 app-bundle, these actions must be able to bring that app-bundle to the
329 foreground.

330 In Apertis, the custom `agl-protocol` protocol extensions are supported, provid-
331 ing the basis to implement notifications and display them as part of the system
332 panel. Currently this feature is not implemented.

333 Focus-stealing

334 A graphical program might attempt to get the user’s attention by creating new
335 main windows while it is in the background.

- 336 • These windows must not be displayed or given input focus, to avoid user
337 distraction and **focus-stealing**¹⁵.
- 338 • We recommend encouraging application developers to use **notifications**
339 instead.
- 340 • Some programs ported from non-Apertis environments might rely on the
341 ability to create a window at any time as a way to get the user’s attention.
342 If a program does this, the compositor must not display it or give it input
343 focus until the user requests **main window switching**.
 - 344 – The compositor could handle this with no user distraction at all, by
345 making the window available in the **main window selection** list, but

¹⁵https://en.wikipedia.org/wiki/Focus_stealing

- 346 not showing it. However, this would not have the desired effect of
347 informing the user that something has happened.
- 348 – Additionally, the compositor could optionally provide a visual cue to
349 the user while minimizing distraction, by behaving as though that
350 program had requested a notification, with content based on the pro-
351 gram and/or window title, and one action button which would bring
352 the new window to the foreground.
 - 353 • If the window would exceed a limit on the number of simultaneous windows
354 or graphical programs in an app-bundle, as described in **main window**
355 **selection**, the compositor must not display those excessive windows, and
356 may terminate the graphical program.

357 As part of the restrictions imposed by **agl-compositor**, only one main window
358 is allowed per application, so the focus-stealing is not impossible. In order to
359 request user attention, applications should use **notifications**.

360 Non-graphical programs

361 A previously non-graphical program could connect to the display server and
362 create a new main window, becoming a graphical program. This situation leads
363 to similar issues as described in **focus-stealing**¹⁶.

364 In order to simplify the design and make it more coherent, applications should be
365 either non-graphical or graphical across all their lifetime. Applications requiring
366 special behaviour should solve this by decoupling their graphical part from the
367 non-graphical by creating a graphical application and a service. In order to use
368 the graphical interface a service can request user attention by using **notifications**

369 Screenshots

370 Screenshots impose a security risk that should be considered.

- 371 • A program from one app-bundle must not be able to copy the texture data
372 of a window from a different app-bundle, which might contain confidential
373 information. (*Output confidentiality*)
 - 374 – In particular, this forbids taking screenshots of a program from a
375 different app-bundle.
 - 376 – The ability for programs in the same app-bundle to take screenshots
377 of each other is optional. For “least-privilege”, we suggest that the
378 platform should not allow app-bundles to request that the platform
379 takes a screenshot of that app-bundle. The programs can communi-
380 cate directly with each other to share their texture data, if desired,
381 so the platform’s involvement is not needed.
- 382 • A program from an app-bundle must not be able to copy the texture data
383 of platform UI elements, which might contain confidential information.
384 (*Output confidentiality*)

¹⁶https://en.wikipedia.org/wiki/Focus_stealing

385 – In particular, this forbids screenshots again.
386 • In some scenarios specially privileged app-bundles must be able to take
387 screenshots, bypassing the restrictions mentioned.
388 • Screencasting or video recording is essentially equivalent to an ongoing
389 stream of screenshots, and has equivalent requirements.

390 On **agl-compositor** this functionality is provided through protocol extensions
391 that can be implemented by a privileged application.

392 **Synthesized input**

- 393 • A program from one app-bundle must not be able to synthesize input
394 events for delivery to a window in a different app-bundle, which could be
395 used to force the target program to carry out undesired actions. (*Input*
396 *integrity*)
- 397 • A program from one app-bundle must not be able to synthesize input
398 events for delivery to the compositor, which could be used to force the
399 compositor or other programs to carry out undesired actions. (*Input in-*
400 *tegrity*)

401 **Trusted input paths**

402 In some situations the platform may need to ask the user for input, in such a way
403 that the user can be confident that their input will in fact go to the platform and
404 not to a potentially malicious app-bundle. One prominent example of a trusted
405 input path is the “Ctrl+Alt+Del to log in” mechanism in Windows operating
406 systems: Windows does not allow ordinary applications to intercept this key
407 sequence, which means that the user can be confident that the resulting login
408 dialog actually belong to Windows, and not an ordinary application that is
409 mimicking it.

410 GNOME uses *system-modal dialogs* for a similar purpose when carrying out
411 platform-related actions like [asking for confirmation](#)¹⁷ of a potentially dangerous
412 system-wide action or when [unlocking access to stored passwords](#)¹⁸.

- 413 • The compositor must be able to request input from the user regardless
414 of any other factors, for example application windows or notifications.
415 (*Availability, integrity*)
- 416 • Other platform components might need to request input from the user in
417 a similar way.
- 418 • Unprivileged app-bundles must not be able to make equivalent requests.
419 (*Output integrity; output availability for everything else*)
- 420 • The trusted input path must be displayed in such a way that only the
421 compositor or another trusted service can produce the same visual result,
422 for example by displaying it over the top of **Platform UI elements** in a

¹⁷<https://wiki.gnome.org/Design/OS/AuthorizationDialog>

¹⁸<https://wiki.gnome.org/Design/OS/KeyringDialog>

423 way that would not be possible or would not be allowed for an ordinary
424 application window. (*Output integrity, input integrity*)

425 As previously commented, **agl-compositor** enforces that applications only cre-
426 ate one main window and standard popup surfaces are not supported. However,
427 using its protocol extensions, a privileged application can display a popup sur-
428 face to provide a trusted input path.

429 Further reading

- 430 • [Wayland Protocol security and authentication](#)¹⁹
- 431 • [Security in Wayland-Based DEs](#)²⁰
- 432 • [Wayland Compositors - Why and How to Handle Privileged Clients](#)²¹
- 433 • [User Interaction Design for Secure Systems \(Ka-ping Yee, 2002\)](#)²²
- 434 • [The status of Wayland security](#)²³

¹⁹<https://wayland.freedesktop.org/docs/html/ch04.html#sect-Protocol-Security-and-Authentication>

²⁰<https://www.x.org/wiki/Events/XDC2014/XDC2014DodierPeresSecurity/xorg-talk.pdf>

²¹<http://www.mupuf.org/blog/2014/02/19/wayland-compositors-why-and-how-to-handle/>

²²<http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.65.5837>

²³<https://lwn.net/Articles/589147/>