

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

 $\mathbf{13}$

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

¹⁹ In Apertis the compositor runs as the agl-compositor¹ 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.

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

We aim to provide the usual security properties described in the Security design $document^2$:

- ³² confidentiality
- integrity
- availability

³⁵ for the two mechanisms provided by the compositor:

• output (placing application windows on the screen)

¹https://docs.automotivelinux.org/en/master/#5_Component_Documentation/1_aglcompositor/

 $^{^{2}} https://www.apertis.org/concepts/archive/application_security/securi$

 input (dispatching input events such as touchscreen touches and gestures to applications)

³⁹ [Wayland Compositors - Why and How to Handle Privileged Clients] provides ⁴⁰ a good overview of how those security properties apply to compositors.

41 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 experi-44 ence (UX) will be a point of differentiation for OEMs, each of these requirements 45 should be interpreted as a requirement that it is possible for the platform to be-46 have as specified, and a recommendation that OEMs'platform variants should 47 do so unless it conflicts with their desired UX. For example, for brevity, we 48 will use "the compositor must ..." as shorthand for "it must be possible for the 49 compositor to ..., and we recommend that OEMs' compositors should have that 50 behaviour unless it conflicts with their desired UX". 51

52 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 point³ results in the relevant graphical program being started.

⁶¹ (These are aspects of input and output availability).

⁶² On Apertis, this job is accomplish by **maynard** which is the default desktop ⁶³ manager.

⁶⁴ 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.

 $^{{}^{3}} https://www.apertis.org/concepts/archive/application_framework/application-entrypoints/$

70	- For example, the equivalent features in Android are the small re-
71	gion 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).

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

⁹² Trusted output

98

99

The compositor must not allow unprivileged programs to display their content in the regions of the screen that are reserved for these UI elements, unless the compositor's UX design specifically allows it. This is a *trusted path* with which the platform can display information to the user. (Output integrity)

- Ideally, the APIs provided to programs should be designed so that it is impossible to request display in a forbidden area.
- If the APIs provided to programs are such that the program can
 attempt to display in these regions, and an unprivileged program
 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

When a graphical program is launched, after carrying some non-graphical initialization, it will create a surface, fill it with the first frame that it wants to be 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 UA involves floating or cascading windows (as seen in CNOME Windows etc.) the manhiel means which had improve the dimensional in
141	in GNOME, windows, etc.), the graphical program must be displayed in the location chosen by the composition. It may influence that location by
142	the location chosen by the compositor. It may influence that location by
143	these hints
144	• The compositor must arrange for any platform III elements that should
145	• The compositor must arrange for any platform of elements that should romain visible at all times to romain visible and interactive during this
140	process (input and output availability).
141	- if they are provided by the compositor itself they must be layered
1/0	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

 $[\]label{eq:approx} $4 https://www.apertis.org/glossary/#app-bundle $5 https://www.apertis.org/glossary/#user-account $2 https://www.apertis.org/glossary/wuser-account $2 https://www.apertis.org/glossary/wuser-account $2 https://www.apertis.org/glossary/wuser-account $2 https://www.apertis.org/glossary/wuser-account $2 https://wwwwwwwwwwwwwwwwwwwwwwwwwwwwwwww$

am

153

- The compositor must deliver location-specific input events such as touchscreen touches to the application at the relevant location, and to no other application. (Input availability, input confidentiality)
- In particular, if application windows can overlap (for example stacking or cascading), and application A is in front of application B, then application A must not be able to trick the user into entering confidential input that was intended for application B by making itself transparent or almost-transparent, so that the user interface of application B shows through (clickjacking⁶). (Input confidentiality)

The compositor must deliver non-location-specific input events such as touchscreen edge-swipe gestures to the current application, using a definition of "current" that is part of its UX, and to no other application. (Input availability, input confidentiality)

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

The Wayland protocol operates via an AF_UNIX socket⁷, just like D-Bus, so applications can be identified by their AppArmor profile and uid using the same credentials-passing mechanisms that are already available in D-Bus.

Also, since user applications are meant to be deployed and launched using the
 Apertis application framework⁸, applying the principles described in security
 boundaries and thread model⁹ minimises the risks.

177 Last-used mode

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

• The platform must arrange for each of the graphical programs that was previously active and visible (in the foreground) to be restarted.

• When one of those graphical programs asks the compositor to display a surface, the compositor must place it in the same location where it was previously visible.

The platform may launch other graphical programs that were running but
 not visible when the state was saved. They must not become visible until

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

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

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

 $^{^{9} \}rm https://www.apertis.org/concepts/archive/application_security/security/#security-boundaries-and-threat-model$

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

to switch to them.

¹⁹² (Input and output availability)

¹⁹³ Main window selection

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

A graphical program might make it difficult for the user to leave, either accidentally (because the program has become unresponsive) or deliberately as a denial
of service (because the program is maliciously written or has been compromised
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

 $[\]frac{10}{10} \rm 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.
	The compositor enforce timeoute when interacting with surfaces in order to
256	The compositor enforce timeouts when interacting with surfaces in order to
257	prevent not responding application to compromise user experience.
258	Child windows
200	
259	A graphical program might include dialogs ¹¹ in its $\bigcup X$.
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

• If surfaces (windows) are thed, stacked of hoating, the dialog may extend outside the boundaries of the graphical program's main window if desired,

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

- but we recommend that this pattern is discouraged. If this is done, it should always be made obvious which surface the dialog belongs to. (Output integrity)
- The dialog must not prevent the user from [leaving the program], even if it extends outside the main window; in other words, it may be app-modal or document-modal, but must not be system-modal. (Input and output availability)
- We suggest encouraging the use of document-modal dialogs 12 similar to those in OS $\rm X^{13}$ and GNOME 14

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

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

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

The support of system popups is implemented through custom protocol extensions which can be only used by privileged applications which implement them, like **maynard** does.

290 Notifications

External events might result in a *notification*, typically implemented as a "pop up"window.

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

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

- If there is a current notification, the platform should draw a visual representation of it, displaying it "above" any current window. (Output availability for the notification)
- If there is no current notification, any program (including non-graphical programs) may trigger a new notification. (Output availability for the notification)

 $^{^{12}} https://en.wikipedia.org/wiki/Dialog_box\#Document\%20modal$

 $[\]label{eq:approx} {}^{13} https://developer.apple.com/library/mac/documentation/UserExperience/Conceptual/OSXHIGuidelines/WindowDialogs.html#//apple_ref/doc/uid/20000957-CH43-SW2$

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

- Each notification should be visually associated with the appropriate appbundle, perhaps via an icon and title. (Output integrity)
- Notifications should be drawn in such a way that only the compositor (or the trusted notification service, if separate) can produce the same visual result, for example by displaying it over the top of platform UI elements in a way that would not be possible or would not be allowed for an ordinary application window. (Output integrity)
- There should be a straightforward mechanism by which the driver can close any notification, minimizing distraction. (Input and output availability for other UI components)
- High-priority platform components such as navigation must be able to force their notifications to be displayed instead of, or "above", other components'notifications. (Output availability for the higher-priority notification)
- Excessive notifications by an application might be distracting. The compositor must have the opportunity to limit the number of notifications per app-bundle or deny notification display altogether, with an optional user-configurable limit per application so that the user could selectively silence an app-bundle that they found distracting.
- The precise handling of notifications (for example topics such as how multiple simultaneous notifications are handled) is outside the scope of this document.
- If the notification has "actions", for example a button to go to the relevant app-bundle, these actions must be able to bring that app-bundle to the foreground.

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

³³³ Focus-stealing

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

- These windows must not be displayed or given input focus, to avoid user distraction and focus-stealing¹⁵.
- We recommend encouraging application developers to use notifications instead.
- Some programs ported from non-Apertis environments might rely on the ability to create a window at any time as a way to get the user's attention.
 If a program does this, the compositor must not display it or give it input focus until the user requests main window switching.
- The compositor could handle this with no user distraction at all, by
- making the window available in the main window selection list, but

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

i	not showing it. However, this would not have the desired effect of
,	informing the user that something has happened.
:	- Additionally, the compositor could optionally provide a visual cue to

the user while minimizing distraction, by behaving as though that program had requested a notification, with content based on the program and/or window title, and one action button which would bring the new window to the foreground.

• If the window would exceed a limit on the number of simultaneous windows or graphical programs in an app-bundle, as described in main window selection, the compositor must not display those excessive windows, and may terminate the graphical program.

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

³⁶⁰ Non-graphical programs

A previously non-graphical program could connect to the display server and create a new main window, becoming a graphical program. This situation leads to similar issues as described in focus-stealing¹⁶.

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

369 Screenshots

374

375

346 347 348

353

354

355

356

370 Screenshots impose a security risk that should be considered.

- A program from one app-bundle must not be able to copy the texture data of a window from a different app-bundle, which might contain confidential information. (Output confidentiality)
 - In particular, this forbids taking screenshots of a program from a different app-bundle.
- The ability for programs in the same app-bundle to take screenshots
 of each other is optional. For "least-privilege", we suggest that the
 platform should not allow app-bundles to request that the platform
 takes a screenshot of that app-bundle. The programs can communicate directly with each other to share their texture data, if desired,
 so the platform's involvement is not needed.
- A program from an app-bundle must not be able to copy the texture data of platform UI elements, which might contain confidential information. (*Output confidentiality*)

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

- In particular, this forbids screenshots again.

- In some scenarios specially privileged app-bundles must be able to take screenshots, bypassing the restrictions mentioned.
- Screencasting or video recording is essentially equivalent to an ongoing stream of screenshots, and has equivalent requirements.

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

³⁹² Synthesized input

385

- A program from one app-bundle must not be able to synthesize input events for delivery to a window in a different app-bundle, which could be used to force the target program to carry out undesired actions. (Input integrity)
- A program from one app-bundle must not be able to synthesize input events for delivery to the compositor, which could be used to force the compositor or other programs to carry out undesired actions. (Input integrity)

401 Trusted input paths

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

GNOME uses system-modal dialogs for a similar purpose when carrying out
platform-related actions like asking for confirmation¹⁷ of a potentially dangerous
system-wide action or when unlocking access to stored passwords¹⁸.

- The compositor must be able to request input from the user regardless of any other factors, for example application windows or notifications. (Availability, integrity)
- Other platform components might need to request input from the user in a similar way.
- Unprivileged app-bundles must not be able to make equivalent requests. (*Output integrity; output availability for everything else*)

• The trusted input path must be displayed in such a way that only the compositor or another trusted service can produce the same visual result, for example by displaying it over the top of Platform UI elements in a

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

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

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

As previously commented, agl-compositor enforces that applications only create one main window and standard popup surfaces are not supported. However,
using its protocol extensions, a privileged application can display a popup surface to provide a trusted input path.

429 Further reading

- Wayland Protocol security and authentication¹⁹
- Security in Wayland-Based DEs²⁰
- Wayland Compositors Why and How to Handle Privileged Clients²¹
- User Interaction Design for Secure Systems (Ka-ping Yee, 2002)²²
- The status of Wayland security²³

 $[\]frac{19}{\rm https://wayland.freedesktop.org/docs/html/ch04.html\#sect-Protocol-Security-and-Authentication}$

 $^{^{20}\}rm https://www.x.org/wiki/Events/XDC2014/XDC2014DodierPeresSecurity/xorg-talk.pdf <math display="inline">^{21}\rm http://www.mupuf.org/blog/2014/02/19/wayland-compositors-why-and-how-to-handle/$

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

²³https://lwn.net/Articles/589147/