



Points of interest

Contents

2	Use-cases	2
3	General points of interest	2
4	TPEG	3
5	Specific points of interest	4
6	Weather	4
7	Security and privacy considerations	4
8	Access to location information	4
9	Attack surface	4
10	Other requirements	5
11	Open questions	6
12	Recommendations	6
13	TPEG	6
14	General POI providers	8
15	Specific POI providers	9
16	Weather	10
17	Dealing with multiple categories	10

Use-cases

General points of interest

Third-party applications (the “provider”) might be aware of general “points of interest” for mapping. For example, an accommodation booking app-bundle might be aware of the locations of hotels, either from a particular chain if the app-bundle is for that chain, or for all hotels if the app-bundle is for a general service like Trip Advisor.

- It must be possible for a third-party app-bundle to provide these “points of interest” to a navigation app.
- It must be possible for the third-party app-bundle to contain a non-GUI service (an “agent”) which will act as the POI provider.
- The navigation app should not be flooded with irrelevant points of interest, for example a hotel in Cambridge while driving from Hannover to Hildesheim
- For privacy, if an application bundle advertises that it is a points-of-interest provider but its manifest indicates that it is not intended to have access to location data, the navigation app must not reveal the current location to it.
- If the POI provider is implemented in terms of an Internet service that can provide a filter/search-based stream of points of interest, it must be possible for the POI provider to avoid being flooded with irrelevant points of interest by the remote server.
- This pattern could potentially be generalized to replace the navigation app with any other consumer.

- We will assume that there is no particular preferred encoding that will be used by a majority of applications, so re-encoding into a common format will usually be necessary.

TPEG

One specific instance of **General points of interest** is receiving **TPEG**¹ broadcasts. These are typically carried on digital radio (DAB) and can encode time- and location-bound events such as traffic congestion, road closures and weather. They can also be transferred via the Internet.

When TPEG is broadcast over DAB, this is done with a “carousel”² approach in which the transmitter repeatedly cycles through a list of currently-valid messages. It is not currently clear to us whether TPEG feeds over the Internet follow the same “carousel” design.

It is not currently clear to us whether TPEG feeds over the Internet filter for relevant events at the client- or server-side.

- It must be possible for an app-bundle to contain an agent that receives TPEG broadcasts and provides events to a navigation app.
- The navigation app should not be flooded with irrelevant points of interest, for example traffic congestion that is not close to either the current location or the projected route.
 - If filtering is performed at the server side, the TPEG provider agent should be able to give the server a suitable filter.
 - If filtering is performed at the client side, it could take place in the agent rather than the navigation app, to minimize IPC traffic between them.
 - If filtering is performed in the navigation app itself, the channel between the TPEG provider agent and the navigation app must be one that is suitable for high-throughput events, and should be able to feed back that events are being transferred faster than the navigation app can read them, so that the TPEG provider agent can “back off” by omitting some events.
- More than one TPEG provider can be installed at the same time. If they are, they are all used in parallel.
- A brochure from the Institut für Rundfunktechnik³ uses 64 kbit/s = 56 messages/s as an indicative figure for TPEG over DAB. This implies that an average binary message should be in the range 100-200 bytes.
- The European Broadcasting Union’s document “TPEG - What is it all about?”⁴ suggests that TPEG-over-Internet clients are expected to poll servers, using a model in which the client downloads a comparatively large

¹<http://tisa.org/technologies/tpg/>

²https://en.wikipedia.org/wiki/Data_and_object_carousel

³https://www.easyway-its.eu/sites/default/files/EW-Highlight_Bavaria_Multimodal_TPEG.pdf

⁴<https://tech.ebu.ch/docs/other/TPEG-what-is-it.pdf>

80 “initial state” at startup, and subsequently receives “deltas” from that initial
81 state.

82 **Specific points of interest**

83 Third-party applications (the “provider”) might be aware of “points of interest”
84 that are relevant to the user with a high probability. For example, an accommo-
85 dation or travel booking app-bundle might be aware that the user has booked
86 a stay at a particular hotel, or a flight from a particular airport; or a PIM ap-
87 plication might be aware of the location of the user’s home or normal place of
88 work, or an upcoming appointment.

- 89 • It must be possible for a third-party app-bundle to provide these “points
90 of interest” to a navigation app.
- 91 • Unlike **General points of interest**, being flooded with irrelevant points of
92 interest is unlikely to be a problem here, because the “provider” application
93 knows that these points of interest are highly likely to be relevant.
- 94 • Time-based filtering could be useful here: appointments beyond a defined
95 date range might not be considered relevant.
- 96 • This pattern could potentially be generalized to replace the navigation
97 app with a non-specific “sink”.

98 **Weather**

99 Weather information might be queried from an Internet service or decoded from
100 **TPEG** broadcasts by an agent in a built-in, preinstalled or third-party app-
101 bundle.

- 102 • The requirements closely resemble **General points of interest**: we require
103 the weather at one or a few weather stations closest to our location or
104 projected route, and we are not interested in distant weather reports.
- 105 • Date-based queries (e.g. projected weather for the next week while plan-
106 ning a long drive) might also be useful here.
- 107 • We anticipate that the data rate here will be considerably lower than for
108 TPEG in general.

109 **Security and privacy considerations**

110 **Access to location information**

111 As noted above, if an application bundle advertises that it is a points-of-interest
112 provider but its manifest indicates that it is not intended to have access to
113 location data, the navigation app must not reveal the current location to it.

114 **Attack surface**

115 TPEG messages come from DAB broadcasts or the Internet.

116 Regardless of how well we might mitigate attacks, if an attacker is able to send
117 crafted TPEG messages to the Apertis system in a way that is indistinguishable
118 from legitimate data, then that attacker can cause the navigation app to display
119 points-of-interest of their choice. This is unavoidable, and is mentioned here only
120 for completeness.

121 Beyond that, if attacker can cause the Apertis system to receive crafted TPEG
122 messages, they might be able to exploit implementation errors in a TPEG parser.
123 We consider three threat-models here:

- 124 • assume that the parser is robust and will not crash or misbehave with
125 malicious input;
- 126 • alternatively, assume that the parser has a denial-of-service vulnerability
127 which causes it to crash or otherwise cease to process information, but
128 does not allow arbitrary code execution;
- 129 • alternatively, assume that the parser has an arbitrary code execution vul-
130 nerability which causes it to execute attacker-chosen code

131 The second and third threat models should be considered to be vulnerabilities
132 (security-sensitive bugs) in the TPEG-parsing component; and can be made less
133 likely by using techniques such as [fuzz testing](#)⁵ to verify the robustness of the
134 parser. However, it might be considered valuable to mitigate any vulnerabilities
135 in those classes that remain and are discovered by an attacker.

136 Similarly, a non-TPEG-based points-of-interest or weather information provider
137 is likely to receive data in some non-TPEG format from the Internet, and equiv-
138 alent security considerations apply to that data. We recommend that crypto-
139 graphically protected channels such as HTTPS are used where possible.

140 Other requirements

141 It is undesirable to increase coupling between consumers and providers by hav-
142 ing consumers provide any location/route information to providers. Instead,
143 providers should retrieve that information from the platform.

144 Similarly, it would be possible for the consumer to give the provider an indication
145 of what is required (in terms of level of detail, search radius around the location
146 and so on). This would allow the provider to optimize its trade-off between
147 resource consumption and information provided, by providing what is needed
148 but no more. However, this causes relatively tight coupling between consumers
149 and providers, which is undesired in an app-centric model. This should not
150 be implemented; instead, the provider should be responsible for determining a
151 reasonable policy.

⁵https://en.wikipedia.org/wiki/Fuzz_testing

152 Open questions

- 153 • Does TPEG already have a defined encoding for points-of-interest similar
154 to those discussed in **General points of interest** and **Specific points of**
155 **interest**, or does it only have encodings for traffic, weather information,
156 and a few specific classes of point-of-interest such as parking?
- 157 • Are the data rates discussed above (64 kbit/sec, 56 messages/sec) repre-
158 sentative?
- 159 • How frequently do we anticipate that a consumer would wish to consume
160 some families of location-sensitive data (“applications” in TPEG jargon -
161 traffic information, weather, points of interest) but not others?
- 162 • Do we expect weather information to come from TPEG, or from a
163 query/lookup-based web service with requests like “tell me the weather in
164 Hildesheim tomorrow”, or a combination of the two?

165 Recommendations

166 See **Data sharing**⁶ for background information on various possible communica-
167 tion between apps. The *consumer* here is the navigation app, and the *providers*
168 are the POI, TPEG and weather providers.

169 We anticipate that the navigation app (or other consumer) might either be a
170 HMI that is started via user action or a platform component that is started
171 automatically on boot, whereas the various providers will be agents provided
172 by platform components and/or store applications, started either on boot or on-
173 demand. Because the consumer might be a HMI, we recommend that it should
174 act as the *initiator* as described in **Data sharing**⁷.

175 We recommend that the navigation app should locate suitable POI, weather
176 and/or TPEG providers by performing **interface discovery**⁸.

177 To bypass concerns about **Access to location information** and ensure that con-
178 sumers and providers remain “loosely coupled”, we recommend that the con-
179 sumer does not inform providers about the current location or route. Instead,
180 each provider that requires this information should retrieve it from a platform
181 service via a D-Bus API, following the conventional publish/subscribe model.
182 The platform should only allow this access for providers whose app-bundle man-
183 ifests authorize it.

184 TPEG

185 For TPEG, there is a choice between several solutions. The answers to the **Open**
186 **questions** above influence the choice between these options.

⁶https://www.apertis.org/architecture/data_sharing/

⁷https://www.apertis.org/architecture/data_sharing/

⁸https://www.apertis.org/concepts/interface_discovery/

187 Our provisional recommendation is to take the **TPEG stream** design, and mini-
188 mize exploitable bugs in TPEG parsing by subjecting the parser to fuzz testing
189 and code auditing.

190 Where TPEG is received with a **carousel**⁹ model, we recommend that the TPEG
191 provider is responsible for keeping a cache of items received during previous
192 sessions, forgetting those items when they are no longer valid or relevant, and
193 replaying them when each new consumer connects.

194 **TPEG stream** One option is to use **data sharing#Consumer-initiated push**
195 **via a stream**¹⁰. In this model, the method that is called by the provider would
196 start a stream of TPEG messages, in either the binary or XML encoding, with
197 a suitable framing protocol (which will need to be specified) if one is required.

198 This solution has the advantage that it allows a TPEG provider to treat TPEG
199 as opaque: for example, a DAB radio receiver that is primarily designed to play
200 audio, but receives TPEG via DAB as a side-effect, could pass TPEG messages
201 through to navigation software without parsing or understanding them. This
202 avoids requiring TPEG parsing in DAB applications or agents, leading to high
203 throughput and low resource consumption by these processes (but potentially
204 a correspondingly higher resource consumption for the consumer, which does
205 need to parse the TPEG).

206 The disadvantage of this solution is that if a provider takes this TPEG pass-
207 through approach, the consumer is directly exposed to potentially malicious data
208 received from the network; for example, a malformed TPEG message might be
209 designed to exploit bugs in the consumer's parser.

210 **Stream of some other format** Another option is to use **data sharing#Consumer-**
211 **initiated push via a stream**¹¹, but require the messages carried by the stream
212 to be in some other format instead of TPEG, for example **KML**¹². This would
213 force the provider to parse TPEG and re-encode it in the other format, which
214 comes with an efficiency cost. It also requires the choice and implementation
215 of the other format; if a new format is designed, requirements-gathering and
216 design will be needed, which could be viewed as a waste of effort that could be
217 avoided by reusing an existing format.

218 All the usual trade-offs between data formats - size efficiency, time efficiency,
219 expressiveness, scope for debugging, and so on - apply to the choice of the other
220 format. GVariant, JSON or XML might make a reasonable basis for a data
221 format, but each of those would require a suitable data representation (schema)
222 to be designed and specified, similar to the way KML is layered above XML.

⁹https://en.wikipedia.org/wiki/Data_and_object_carousel

¹⁰https://www.apertis.org/architecture/data_sharing/#consumer-initiated-push-via-a-stream

¹¹https://www.apertis.org/architecture/data_sharing/#consumer-initiated-push-via-a-stream

¹²https://en.wikipedia.org/wiki/Keyhole_Markup_Language

223 If the TPEG parser is assumed to be robust, this approach has no advantage
224 over the TPEG stream. The advantage of this option over a TPEG stream is
225 that if the TPEG parser is not considered to be robust, this option somewhat
226 reduces the **attack surface** available to potentially-malicious TPEG messages.

227 If the attacker is only assumed to be able to cause a crash via a malformed
228 message, this is entirely mitigated by performing TPEG parsing in the provider:
229 the consumer would stop receiving messages from that provider, but continue
230 to run. Optionally, it could detect the unexpected end-of-stream and re-initiate
231 communication with a new instance of the provider.

232 However, if the attacker is assumed to be able to cause arbitrary code execu-
233 tion in the TPEG parser, this design does not necessarily provide any mitiga-
234 tion: having subverted the TPEG provider, the attacker could send arbitrary
235 messages to the consumer in the other format, potentially triggering denial-
236 of-service or code-execution vulnerabilities in the parser for that other format.
237 This could be mitigated by requiring that the parser for the other format has a
238 robust and well-understood implementation provided by the platform.

239 **Publish/subscribe via D-Bus** TPEG providers could use a **pub-**
240 **lish/subscribe approach via D-Bus**¹³, with each D-Bus message either
241 carrying a binary or XML TPEG message, or a message translated into a
242 non-TPEG encoding using D-Bus data structures. However, our current
243 understanding of the expected message rate is that it is at the high end of the
244 rates for which D-Bus was designed, so we do not recommend this.

245 In addition, if the TPEG provider does not cache currently-valid TPEG mes-
246 sages in memory but merely passes them through as they are received, then
247 this would be a poor fit for conventional D-Bus design patterns, since a `GetCurrentState()`
248 method call would not make sense (the provider would not have any
249 current state to get).

250 **General POI providers**

251 For POI providers that are implemented in terms of access to a web service,
252 either the POI provider can perform a series of search operations on the remote
253 server and get relevant POIs in response (a “pull”model), or the remote server
254 can send a stream of POIs to the POI provider while periodically polling the
255 current locations of interest (a “push”model).

256 For the transport between the POI provider and the POI consumer, multiple
257 approaches are possible, depending on the answers to the **Open questions**.

258 Our provisional recommendation is that points of interest are encoded as TPEG
259 and transferred to consumers via the same mechanisms we would use for a TPEG
260 stream.

¹³https://www.apertis.org/architecture/data_sharing/#publish.2fsubscribe-via-d-bus

261 **TPEG stream** If TPEG has an encoding for points of interest, then the
262 same approach can be taken as for TPEG: [data sharing#Consumer-initiated](#)
263 [push via a stream](#)¹⁴, with a stream of binary or XML TPEG messages. This
264 has the advantage that we do not need to define our own encoding for points of
265 interest.

266 Another advantage of this approach (assuming that TPEG providers also pro-
267 vide a TPEG stream) is that the consumer only needs to consume TPEG, rather
268 than consuming both TPEG (from DAB broadcasts) and some other format
269 (from general POI providers). If TPEG does not already have an encoding for
270 points of interest, but it is feasible to add one, it might still be worthwhile to
271 do so in order to receive the second advantage.

272 The disadvantage is that each provider needs to encode whatever information it
273 provides into TPEG format. This could introduce a loss of information if TPEG
274 is not sufficiently expressive to describe everything the provider requires.

275 If the corresponding solution is not chosen for TPEG providers, this approach
276 should not be chosen either.

277 **Stream of some other format** As with TPEG, we could define a non-TPEG
278 encoding for points of interest and stream them to the consumer using [Data](#)
279 [sharing#Consumer-initiated push via a stream](#)¹⁵. Similar considerations apply
280 to the design of the new format: we could select an existing format such as
281 [KML](#)¹⁶, or design a new one, perhaps using GVariant or JSON as a base.

282 If the non-TPEG encoding is chosen well, it should be possible to encode all
283 aspects of a point of interest without information loss. However, if the non-
284 TPEG encoding is not sufficiently expressive, information might still be lost in
285 translation.

286 **Publish/subscribe via D-Bus** As with TPEG, POI providers could use a
287 [publish/subscribe approach via D-Bus](#)¹⁷, with each D-Bus message carrying one
288 or more points of interest, either described using D-Bus data structures or as
289 an opaque byte-array in some known format. We anticipate that POI providers
290 would normally have rather lower message rates than TPEG, so they might be
291 within the range for which D-Bus is designed, even if TPEG is not.

292 **Specific POI providers**

293 The implementation options for specific POI providers are essentially the same
294 as general POI providers, although their requirements are less stringent: the

¹⁴https://www.apertis.org/architecture/data_sharing/#consumer-initiated-push-via-a-stream

¹⁵https://www.apertis.org/architecture/data_sharing/#consumer-initiated-push-via-a-stream

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

¹⁷https://www.apertis.org/architecture/data_sharing/#publish.2fsubscribe-via-d-bus

295 message rate is anticipated to be lower, and location information is not neces-
296 sarily needed.

297 We recommend that the same interface as for general points of interest is used:
298 we do not see any reason why we need to distinguish between the two.

299 Weather

300 One possible implementation is to use the same TPEG-stream-based design as
301 the other use cases. This is appropriate if we anticipate that the majority of
302 providers of weather information will be implemented using TPEG or similar,
303 but it is a poor fit for a more query-based provider such as one that accesses
304 a web service: the provider would have to perform speculative queries covering
305 the area around the location and/or route, and emit their results at intervals as
306 TPEG or TPEG-like messages.

307 Another possible implementation is to use [Data_sharing#Query-based access](#)
308 [via D-Bus](#)¹⁸, which is a good fit for implementations that query a web service,
309 but a poor fit for implementations that receive weather from TPEG (which
310 would have to cache all available weather information, and use their cache to be
311 able to answer queries). This could be mitigated by treating weather queries as
312 a platform service, so that the cache is maintained by the platform.

313 A third possibility is to have two separate weather APIs - one stream-based de-
314 sign for “ambient weather information”, and one query-based design for “current
315 and forecasted weather queries”- and require consumers to choose whether to
316 use one or both depending on their particular requirements. This design has the
317 advantage that each of those APIs represents the underlying weather service in
318 the most natural way, but has the disadvantage that consumers are expected
319 to use two parallel APIs.

320 We do not currently have a specific recommendation here, and would appreciate
321 feedback on the relative priorities of those designs’ strengths and weaknesses.

322 Dealing with multiple categories

323 This recommendation is conditional on the answers to the [Open questions](#),
324 above.

325 We recommend using a name such as “categories”, “types” or “classes” for TPEG’s
326 jargon term “applications”, to avoid confusion with the already overloaded term
327 “application”. Documentation could mention the TPEG jargon for clarification,
328 for example:

```
329 /** * ... * Return a list of categories of location information that are sup-  
330 ported * by this provider. These categories are the same concept as * "applica-  
331 tions" in TPEG terminology. * ... */ gchar **..._list_categories (... *self);
```

¹⁸https://www.apertis.org/architecture/data_sharing/#query-based-access-via-d-bus

332 If we anticipate that all categories of data will typically all have the same con-
333 sumers, we recommend having a single shared interface for [interface discovery](#)¹⁹,
334 perhaps `org.apertis.LocationInfoProvider` or `org.apertis.TPEGProvider`. Clients
335 not requiring all of the available data for that interface could receive and discard
336 it.

337 If we anticipate that categories will often have different consumers, we will need
338 a list of categories, and an interface for [interface discovery](#)²⁰ per category, for
339 example `PointsOfInterestProvider`, `TrafficProvider` and `WeatherProvider`.

340 One approach to these categories would be to define a separate D-Bus interface
341 per interface-discovery interface, with intentionally similar D-Bus APIs to set
342 up a stream. Each provider that could provide more than one category would
343 be required to demultiplex the messages that it receives and write them into
344 the appropriate streams.

345 Another approach to these categories would be to define a separate D-Bus in-
346 terface per interface-discovery interface, with intentionally similar D-Bus APIs
347 to set up a stream. Each provider that could provide more than one category
348 would be required to demultiplex the messages that it receives and write them
349 into the appropriate streams.

¹⁹https://www.apertis.org/concepts/interface_discovery/

²⁰https://www.apertis.org/concepts/interface_discovery/