

1

Federated Music Player

2

Jade Keira Melody Ellis

`jkme2@kent.ac.uk`



School of Computing

University of Kent

3

June 15, 2026

Abstract

5 As streaming services dominate the modern music consumption landscape,
6 listeners often find themselves constrained by heavily limited local music
7 players and centralized platforms with limited licensed libraries that inad-
8 equately compensate artists. This report introduces a project codenamed
9 PlayerBrainz, a federated music library server that aims to bridge the gap
10 between local music collections and the rich metadata available through the
11 MetaBrainz ecosystem. By storing music libraries locally and establishing a
12 robust federation mechanism, users can connect their independently hosted
13 servers, share access to their libraries, and discover new music collaboratively
14 without relying on an intermediary streaming provider. This report details
15 the design and implementation of a scalable Rust back-end and correspond-
16 ing GraphQL API capable of distributed federation, paired with a SvelteKit
17 administrative front-end.

18 **1 Introduction**

19 The transition from physical media and locally managed music collections
20 to centralized streaming platforms has fundamentally altered how users dis-
21 cover, consume, and manage music. While platforms like Spotify offer unpar-
22 alleled convenience and expansive libraries, they introduce rigid constraints.
23 Users sacrifice ownership of their music files (with tracks frequently disap-
24 pearing from user’s libraries as licensing changes), and are subject to recom-
25 mendation systems driven by platform profitability rather than pure artistic
26 alignment. Furthermore, centralized streaming largely disconnects listeners
27 from the underlying creative ecosystems, commoditising music by divorcing
28 it from its context, authorship, and performers.

29 The MetaBrainz Foundation, through its interconnected databases like
30 MusicBrainz and ListenBrainz, has chronicled detailed and intricate rela-
31 tionships spanning millions of musical acts globally. However, modern music
32 players predominantly utilise a naïve schema consisting solely of simple tex-
33 tual fields like ”Title”, ”Artist”, and ”Album”, which frequently run into lim-
34 itations displaying real information, omitting the rich collaborative tapestry
35 of the music industry.

36 To counter these limitations, this project conceptualises and implements
37 PlayerBrainz, a federated infrastructure designed to expose the full depth
38 of MusicBrainz relationships directly alongside the user’s audio files. The
39 vision extends beyond a single local media player; it proposes a distributed,
40 relationship-driven alternative to commercial streaming services. By allowing
41 users to host their music on a private server and federating these servers using
42 standardized cryptographic protocols, individuals can share their libraries
43 with family and close friends, retain absolute ownership of their files, and
44 explore the rich, multifaceted universe of music organically.

45 This report documents the architectural design, technological choices,
46 and the practical implementation of the PlayerBrainz backend and frontend
47 systems.

48 **2 Previous Work**

49 **2.1 The MetaBrainz Foundation**

50 The MetaBrainz Foundation is a non-profit organisation that maintains a
51 number of complementary projects related to music metadata. These were
52 deemed of particular relevance to the project, and are described in more
53 detail below.

54 **2.1.1 MusicBrainz**

55 MusicBrainz (Swartz, 2002) is a collaborative structured database of music
56 metadata. It is the most comprehensive and widely used music metadata
57 database, and is the basis for many other projects in the music metadata
58 space.

59 The MusicBrainz database schema (The MusicBrainz Authors, 2025) pro-
60 vides a structured and accurate representation of music metadata as encoun-
61 tered in the real world. Naïve approaches to music metadata often assume
62 a simple schema where each track has a single title, artist and album, and
63 that albums have a single artist - each of these would be keyed by its hu-
64 man readable name. In reality, music metadata is much more complex. In
65 contemporary music, it is common for tracks to have multiple artists, and
66 for the performing name of artists to change over time. Credits are often
67 done in unique formats. In addition to this, the same track may be re-
68 leased on multiple albums or singles, with multiple different recordings. The
69 MusicBrainz schema accounts for this complexity and provides a structured
70 way to represent it. It also provides a way to link to external databases,
71 including streaming providers. For this reason, it was chosen as the basis for
72 the schema used in this project, with some adjustments to account for the
73 particular use case of on-disk music files.

74 **2.1.2 MusicBrainz Picard**

75 MusicBrainz Picard is an application for tagging and organising audio files
76 using metadata from the MusicBrainz database. It uses multiple techniques,
77 including audio fingerprinting techniques, to match audio files to correspond-
78 ing entities in the MusicBrainz database, and then adds metadata to them ac-
79 cording to configured rules. The project maintains a schema of tag mappings
80 (The Picard Authors, 2026), which can be used by third party programs. No-
81 tably, Picard embeds the MusicBrainz identifiers (MBIDs) of several entities
82 by default.

83 **2.2 Music Servers**

84 Other programs exist which provide the basic functionality of exposing a mu-
85 sic library over the network with an API for the purpose of streaming. How-
86 ever, most use somewhat naïve approaches to music metadata, with many of
87 the shortcomings listed earlier. Additionally, only Funkwhale (Berriot, 2017)
88 has any kind of federation functionality - however, it has a particular focus
89 on publicizing uploaded music using the ActivityPub social protocol.

90 **2.2.1 Subsonic**

91 These music servers have a *de facto* common API known as the Subsonic
92 API (The Subsonic Authors, 2005). This is a JSON and XML API provid-
93 ing a wide variety of related endpoints and functionality for music, podcast
94 and other audio streaming. The OpenSubsonic project (The OpenSubsonix
95 Authors, 2023) is an attempt to formalise this to deal with some of the am-
96 biguities and limitations of the Subsonic API. However, it was decided that
97 avoiding implementing this API would provide greater flexibility to experi-
98 ment with this project, and it was not necessary for the goals of the project.

99 2.3 Federated Protocols

100 In modern federated software, there are two primary relevant protocols – the
101 Matrix protocol (The Matrix Authors, 2026), and the ActivityPub protocol
102 (Lemmer-Webber et al., 2018). Both are comprehensive standards including
103 key parts such as server discovery, HTTP signatures, user identifiers and
104 more to provide everything needed to implement software in their respective
105 ecosystems. The Matrix protocol has a focus on synchronizing state between
106 untrusted servers, with byzantine fault tolerant properties. The ActivityPub
107 protocol focuses on subscription & replication of structured data from the
108 originating servers.

109 Given that the application is intended to synchronize data between trusted,
110 cooperating systems, it was decided that neither protocol could be used in
111 its existing form.

112 3 System Overview

113 3.1 Back-end

114 The server is required to return large amounts of deeply nested data in ar-
115 bitrary structures. For example, a music library page might need to display
116 various tracks, albums, releases, and other related metadata. Because each
117 potential screen design needs to display different subsets of data (e.g., some
118 views require the track year or album while others do not), and because ide-
119 ally third party apps would be able to use the API, it would be impossible
120 to design traditional REST endpoints that deeply returns all necessary data
121 without also returning a significant amount of unnecessary data. Conversely,
122 relying on smaller, limited endpoints would necessitate a large number of
123 API requests to populate a single page, resulting in a slow and poor user
124 experience.

125 GraphQL was selected to solve these problems. It allows clients to query
126 exactly the data they need in one request, including deeply nested struc-

127 tures. When examining the Rust ecosystem for GraphQL implementations,
128 `async-graphql` appeared to be the most popular and established asyn-
129 chronous library. It provided all the necessary features and functionality
130 for the application and was therefore chosen.

131 For database interactions, SeaORM was selected. It is the leading asyn-
132 native Rust object-relational mapper (ORM) and is built upon SQLx, which
133 is the leading asynchronous SQL database client library for Rust.

134 **3.2 Front-end**

135 To demonstrate the API provided by the back-end, a front-end application
136 was necessary. The SvelteKit meta-framework was chosen based on existing
137 developer familiarity and the known positive performance properties of its
138 implementation. The user interface was built without the use of a CSS
139 framework, relying on handwritten basic styles.

140 To integrate with the back-end API, a SvelteKit-compatible GraphQL
141 client library named Houdini was selected. Initially, Houdini provided every-
142 thing necessary, featuring ergonomic code generation and context drilling.
143 However, a significant limitation was discovered later in the project: when
144 Houdini is used with SvelteKit, it relies on one global Houdini instance. As
145 a result, it can only connect to a single GraphQL endpoint at a time, mak-
146 ing logging into and querying arbitrary selfhosted servers significantly more
147 difficult.

148 **3.3 Federation**

149 When making federated requests between instances, a robust way of authen-
150 ticating the server making the request is needed, in addition to the standard
151 transport-layer security (TLS) authentication provided for the response. Mu-
152 tual TLS was considered but was deemed impractical due to a lack of broader
153 ecosystem support and the difficulty of acquiring certificates issued by the
154 existing public Certificate Authority infrastructure.

155 Other request signature schemes were examined, leading to the decision
156 to use RFC 9421 HTTP Signatures. This standard is well-specified, and
157 an existing Rust implementation, the `httpsig` library, was available to be
158 adjusted and integrated into the application. Using a standardised RFC
159 also provides potential future interoperability benefits, even if those are not
160 immediately visible at this stage of the project. The `cynic` GraphQL library
161 is used as the client in the back-end to execute these federated queries, and
162 the application utilises a ‘.well-known‘ scheme for endpoint discovery.

163 4 Implementation

164 4.1 Schema

165 The MusicBrainz schema provided a strong foundation for the application.
166 However, it had a few limitations. Primarily, the MusicBrainz Recording
167 entity was not granular enough to match an actual on-disk recording, because
168 a Recording would be the same regardless of encoding, bitrate, or starting
169 periods of silence. It was decided that a new entity called a ‘Mastering’
170 would be created, with each instance of a Mastering being sample-for-sample
171 identical. As a simplification, each file on disk would be its own unique
172 Mastering.

173 In addition to this, some adjustments to the positioning of cover art and
174 lyrics were made. Cover art corresponded to a release, however on-disk music
175 files can additionally have media embedded within the file, which must be
176 represented. Lyrics are not present in the MusicBrainz schema, and although
177 the Work entity is technically the most correct place to store them, they are
178 attached to a Mastering in typical filesystem representations - even ignoring
179 that not all Recordings have had corresponding Works created and linked.

180 4.2 Scanner

181 Producing a robust and efficient scanner was a particular point of difficulty
182 and iteration. The scanner was originally prototyped as a routine that re-
183 cursively walked the music directory, upserting music files as it found them.
184 However, this handled two important cases poorly. It could not handle mu-
185 sic files being inserted or updated during the lifetime of the server, and it
186 could not handle music files being deleted. For this reason, a new design was
187 created. It also had a number of other limitations, like only being able to
188 handle a single on-disk library.

189 Particular care was taken regarding filesystem watching. It is known that
190 watching large numbers of files is not reliable (Lar, 2022), and the server may
191 be offline for periods of time. Therefore, an incremental update scheme was
192 chosen that could be run from scratch, and was resilient to file watch events
193 being incomplete or incorrect.

194 The new design comprised of two parts. The first part was the file reader
195 process. This process first recursively scanned each directory it was config-
196 ured to watch. It would simultaneously initialize a filesystem watcher, which
197 would start processing events once the recursive scan was done. Directories
198 would be walked in a depth-first manner, sending first the items found in
199 a directory down a channel, and then upon the completion of a particular
200 directory an event marking that directory as complete. The watcher would
201 then mark particular directories as modified and re-scan the subdirectory.

202 The second part was a state machine that received events from the reader
203 process. Relying on the precondition that the channel was a serial depth-
204 first search of the music file, it would upsert the Mastering entities into the
205 database and collect other data. Upon receiving the directory complete noti-
206 fication, it would remove entities that had not been sent, with the assumption
207 that they had been deleted, and finalize the transaction with any other data.

208 Of some note were the issues regarding extracting the tags from the
209 metadata. The upstream library used for audio container parsing seemed
210 to have inconsistent support for metadata across file types, with some man-

211 ual parsing required to consistently extract recording IDs from mp3 contain-
212 ers. Additionally, some tags appeared to be mapped incorrectly, for example
213 `MusicBrainzRecordingId` never appeared in real files, while `MusicBrainzTrackId`
214 contained the actual recording ID. For this reason, the tag extraction code
215 is somewhat more complex than it would otherwise be, and may need to be
216 adjusted in the future if the upstream library is updated.

217 4.3 API

218 The project exposed a discoverable and introspectable API using the GraphQL
219 protocol and conventions. This API provided basic user management, library
220 management and federation functionality. This API is unified between unau-
221 thenticated, authenticated local access, and authenticated federated access.
222 Each subsection of the graph had access control configured as appropriate.
223 Local user authentication followed a simple session token scheme, with de-
224 fined expiration.

225 4.4 Federation

226 The federation protocol took some superficial inspiration from the Matrix
227 protocol, exposing files at well-known endpoints to direct federated servers
228 to the correct API endpoint. Remote servers then make an unauthenticated
229 request to that endpoint to retrieve one or more ed25519 keys associated
230 with the server. Once those keys are retrieved and stored, they may then be
231 used to authenticate requests from that server. However, this was where the
232 similarities ended. The signature scheme used was an RFC 9421 signature
233 of the request with a sha512 digest of the body, rather than a custom HTTP
234 signature scheme, and the federation API was exposed over a GraphQL API
235 described earlier rather than a RESTful JSON API.

236 Requests to the API are authenticated using signed requests. The server
237 fetches the public keys of the requesting server by first requesting the JSON
238 file under the `/.well-known/playerbrainz/server` path on the host. That

239 responds with a JSON object containing a URL under the `graph_endpoint`
240 property, from which it can request the appropriate keys using the GraphQL
241 protocol. Using these keys, the request is then checked, and the server can
242 be authenticated. If the server is authorized to access an endpoint, it can
243 then be used.

244 **5 Results**

245 The project successfully delivered a robust foundation for a federated music
246 library system, alongside a functional administrative front-end. Significant
247 focus was placed on the core systems, establishing a strong base for future
248 development.

249 **5.1 Backend**

250 The Rust-based back-end server successfully implements a GraphQL API
251 leveraging existing ecosystem libraries where appropriate.

252 A prominent achievement of the backend is the implementation of the
253 resilient file scanner. The scanner reliably reads the local filesystem, extracts
254 advanced metadata conforming to the adjusted MusicBrainz schema, and
255 tracks incremental changes.

256 **5.2 Federation**

257 Federation functionality was implemented according to the designed archi-
258 tecture. The system successfully exposes endpoints for node discovery via
259 the `/.well-known/playerbrainz/server` path. Remote servers are able to
260 request necessary public keys, and subsequent cross-instance queries are se-
261 curely verified using RFC 9421 HTTP Signatures. This allows authenticated
262 exchange of data between cooperating servers, enabling the core vision of a
263 federated music library ecosystem.

264 5.3 Front-end Administration

265 The SvelteKit front-end provides the basic administrative user surface for
266 the server, using Houdini to drive state management and GraphQL API
267 connectivity. This includes authentication and basic CRUD functionality for
268 managing the users and libraries on the server.

269 While front-end prototypes for actual music playback and track listing
270 were explored (visible in isolated components of the codebase), they were
271 deferred in the final artifact in favour of establishing a solid federation and
272 API core. This is further discussed in the Future Work section.

273 6 Future Work

274 6.1 MusicBrainz Enrichment

275 One strongly considered functionality of the project was the ability to en-
276 hance the API with metadata retrieved from the MusicBrainz database. This
277 would have allowed an integrated and accessible way of using that informa-
278 tion in downstream apps.

279 Due to limitations and strict rate limits in the MusicBrainz API, func-
280 tionality to retrieve & synchronise metadata with the MusicBrainz database
281 was not developed. A new MusicBrainz API that would relieve these limita-
282 tions is in development as a part of Google Summer of Code 2026, but would
283 not be completed by the conclusion of this project. Using a local mirror of
284 the database was considered as an alternative solution, but it was decided
285 that it would be impractical due to the large download size, and it would be
286 better to wait for other work to mature.

287 6.2 Unified library API

288 During development, an API over the local library was exposed. However,
289 this was removed as federation functionality was developed as it only ade-
290 quately exposed local, already-downloaded music files. Developing an API

291 that consistently exposes remote and local media would make the project
292 more appealing to developers and end users.

293 **6.3 ListenBrainz Integration and Recommendations**

294 A significant future goal for PlayerBrainz would be to more tightly integrate
295 with the ListenBrainz project. ListenBrainz provides a rich source of user
296 listening history, loved/hated tracks, and social feeds. By natively importing
297 this data, PlayerBrainz could offer personalized recommendations and social
298 features that enhance the user experience and increase music discoverabil-
299 ity. Further on from that, using learnings from the Troi recommendation
300 framework, PlayerBrainz could implement music recommendation and radio
301 features that provide functionality familiar to users from commercial stream-
302 ing services.

303 **7 Acknowledgment**

304 Robert Kaye, who never got to see the end of this project. Dan Bard, who
305 has been a great support throughout this. The rest of the MetaBrainz team,
306 who have patiently put up with my nonsense.

307 **References**

- 308 (2022). Large scale watching of paths appears to drop/miss events · Issue
309 #412 · notify-rs/notify.
- 310 Berriot, A. (2017). Funkwhale.
- 311 Lemmer-Webber, C., Jessica, T., Shepherd, E., Guy, A. and Prodromou, E.
312 (2018). ActivityPub.
- 313 Swartz, A. (2002). MusicBrainz: A semantic Web service. *IEEE Intelligent*
314 *Systems*, 17(1), pp. 76–77.

- 315** The Matrix Authors (2026). Matrix Specification version v1.18.
- 316** The MusicBrainz Authors (2025). MusicBrainz Database Schema.
317 https://wiki.musicbrainz.org/MusicBrainz_Database/Schema.
- 318** The OpenSubsonic Authors (2023). OpenSubsonic.
319 <https://opensubsonic.netlify.app/>.
- 320** The Picard Authors (2026). Appendix A: Tag Mapping. [https://picard-](https://picard-docs.musicbrainz.org/en/latest/appendices/tag_mapping.html)
321 [docs.musicbrainz.org/en/latest/appendices/tag_mapping.html](https://picard-docs.musicbrainz.org/en/latest/appendices/tag_mapping.html).
- 322** The Subsonic Authors (2005). Subsonic.
323 <https://www.subsonic.org/pages/api.jsp>.