

Radio Authentication Workshop

23rd August 2012

Attendees

Chris Lewis (BBC)
Sean O'Halpin (BBC)
Chris Needham (BBC)
Dan Nuttall (BBC)
James Harrison (BBC)
Robin Cooksey (Frontier Silicon)
Mathias Coinchon (EBU)
Michael Barroco (EBU)
Tobias Schlatter (EBU)

Agenda

- Overview of RadioTAG specification as it stands (SOH)
- An overview of MediaAUTH (EBU)
- Overview of recent BBC work on authentication for TVs (CN)
- Issues with OAuth2 and its suitability for Radio authentication (CN)
- Applications for authenticated radios (Discussion)

Overview of existing RadioTAG spec

Sean circulated a print out of draft 4 of the RadioTAG 1.00 specification and took us through a presentation. This is the most up-to-date version of the specification as currently implemented on the Revo Axis and BBC prototype radiotag service.

The draft is available here:

<http://radiotag.prototyping.bbc.co.uk/docs/radiotag-api-proposal-v1.00d4.html>

The presentation is available here:

<http://radiotag.prototyping.bbc.co.uk/docs/presentation/radiotag-auth.html>

We discussed the principles behind RadioTAG - specifically the use of tokens to allow restricted access and revocation of credentials.

Sean walked through the unpaired-to-paired narrative.

Sean mentioned the user testing of RadioTAG authentication and that it was generally well-received. We didn't have any complaints from the testing panel about the process of registering the device.

Robin pointed out that he hasn't seen any published terms of service for RadioDNS applications.

We discussed whether the RadioTAG specification allowed device manufacturers to set up "portals" for tagging-type applications. Sean explained that the protocol doesn't preclude this as the manufacturer can set up a proxy service which passes on the tag request to the broadcaster.

Sean listed some issues for consideration

- Business models and a per-broadcast vs centralized service
- What is the minimum device profile we need to support?
- Do we want to use OAuth2 - what is the future of that spec
- Where does identity reside, how does this relate to user switching on the radio
- Are we asking too much of the token in RadioTAG - it does authentication, authorisation - does it do too much?

Tobias has some questions about the technical details of the unpaired grant, but we agree to discuss this later.

Mathias asked what is implemented at the moment. The RadioTAG v4 spec is implemented on the Revo Axis test build and on the BBC's prototype service.

Before v5 Frontier Silicon and the BBC want to implement the SSL support.

Overview of MediaAUTH spec

Michael gave a presentation of the MediaAUTH spec.

The MediaAUTH proposal aims to make it easier for listeners to switch stations, use authenticated services without having to register for a new account.

We had a long discussion around the utility of the proposed directory service and how a device knows which auth providers to trust.

[MediaAUTH presentation and proposal draft](#)

TV Authentication

Chris Needham presented an overview of the BBC's work on authenticating internet-connected TVs.

Chris showed the interface for using a RadioTAG-like flow for authenticating on a TV. Entering the PIN number is quite simple using the TV remote control's numbers.

He then presented an alternative OAuth2-based procedure for authenticating.

We discussed some methods of setting up an authenticating devices where their network address is not known, for example when a mobile phone needs to act as a remote control for the set top box. In this case there is no fixed URL for the set top box, so we are unsure how to meet the requirements set out in the OAuth2 specification with regard to the `redirect_uri` parameter.

Robin mentioned that in some Frontier Silicon devices the device itself acts as a wireless access point that a user can connect to in order to configure the device for the first time. They use this protocol when [WPS](#) is not available.

We discussed OAuth2 and our general opinions. Sean referred to Eran Hammer's [blog post](#) citing his reasons for giving up the editorship of the specification. We discussed the problems with inter-operability. Robin suggested that we could specify which flows we support in our authentication protocols in order to lock down interoperability. Tobias said that OAuth2 is better considered a framework with some design patterns that we can use to ensure security and interoperability for our use cases.

More opinions on OAuth2: <http://gist.io/3170829>. [SaferWeb: OAuth2.a or Let's Just Fix It](#)

Discussion / AOB

Is the future mobile apps that connect to "dumb" speakers? (CRL)

RadioTAG is fundamentally a simple protocol for relating time and broadcast data to information -
can be used on "real" radios, mobile apps, web pages etc. (SOH)

Syndication of tags is a powerful concept (aggregation of tags from different broadcasters) -
what a tag is needs to be defined for this to be possible (RC)

When you control the stack you can use different pieces of the protocols. We should make the distinction between auth and tag more clear in the RadioTAG spec, for example as actually “tagging” is quite a simple protocol which could be used in, for example, a broadcasters mobile app (RC)

Tobias asked some questions about headers in the RadioTAG spec and why we don't use pre-existing authentication headers.

ACTION: Tobias to raise this as an issue on the RadioTAG list.

Mathias asked where to find out about developments to RadioTAG. Robin clarified that the initial draft of the RadioTAG spec is to be published on the mailing list by the editor (Andy Buckingham) and that will be the future point of contact for all discussions.

We discussed the extensibility of the ATOM XML used in the response to tag requests. It would be useful if people could pass data in the ATOM feed for custom clients and so on.

ACTION: Sean to add to the spec that a client shouldn't reject as invalid ATOM that contains extensions that the client doesn't understand

Radio Synchronisation Workshop

24th August 2012

Attendees

Sean O'Halpin (BBC)
Chris Needham (BBC)
James Harrison (BBC)
Robin Cooksey (Frontier Silicon)
Mathias Coinchon (EBU)
Michael Barroco (EBU)
Tobias Schlatter (EBU)

Agenda

- Objectives and use cases (CN)
- Overview of the problem space (SOH)
- Possible solutions (CN)
- Time signalling in broadcast platforms (JH)
- Synchronisation of events to HTML5 video (CN)
- Review of RadioDNS draft document: "Synchronisation of RadioDNS Applications to Audio"
- Discussion
- AOB

Objectives and use cases

Chris asked the group to identify use cases that need synchronisation. We reviewed the list of use cases in the "Synchronisation of RadioDNS Applications to Audio" document and added some new ones:

- Synchronisation of RadioVIS slides and text messages to the audio stream
- Resume listening of a live stream in an on-demand
- Service following: switching between bearers whilst maintaining as closely as possible the same point in the audio
- Personalised schedule, e.g., the ability to switch to an alternate audio stream at a specific point in time in a broadcast

Chris suggested that the use cases can be grouped into several categories:

1. Those that involve the receiver displaying associated content time-aligned with the audio, e.g., VIS slides, text messages, captions, programme chapters/segments
2. Those that involve user feedback, e.g., tagging a song, trail, or advert, or voting
3. Those that involve synchronisation of multiple simultaneous audio or video streams, e.g., delivery of a surround sound audio stream or video. Service following may fall into this category, if it is to be seamless.

The group discussed timing requirements to support these use cases. Category (3) requires sub-millisecond accuracy in order to properly align multiple audio or video streams.

In category (2) we should consider how quickly a listener is likely respond to an item of interest or a call to action. For example, the listener must reach for their radio, remote control, or mobile phone in order to tag a song they like, by which time the next song may have started playing. Consequently it may be difficult for broadcasters to interpret the intention of a listener's tag. (Additional metadata may help with this, e.g., present a list of the last 3 songs and allow the user to choose which one they want to tag.) In any case, in category (2) we want to record the time of tagging as accurately as possible with respect to the audio stream, regardless of the time taken for the user to respond.

Categories (1) and (2) therefore have similar timing requirements. Robin suggested that the resolution on the order of milliseconds is needed, with accuracy in the 250ms to 500ms region, and Chris suggested that 250ms should be the upper limit, if achievable in practice.

James pointed out that the timing resolution sets a limit on the speed at which updates could be expected to be correctly handled by receivers.

Matthias mentioned that listeners complain about the different delays on various bearers or different receivers if they have several radios switched on at the same time in the home.

The group discussed category (3) and came to the conclusion that because of the much higher accuracy requirement this category should be outside the scope of an initial synchronisation specification for RadioDNS applications, in the interest of arriving at something achievable for categories (1) and (2). The group decided not to further consider category (3) at this meeting.

Overview of the problem space

Sean described different methods that can be used to synchronise events to audio. These included embedding timestamps in the media stream, audio fingerprinting, and synchronisation of clocks over the Internet. (Notes here: <http://radiotag.prototyping.bbc.co.uk/docs/presentation/problem-space.html>).

Embedded timestamps

The embedded timestamps method involves delivering a timestamp from the broadcaster to the

receiver, either encoded within the audio signal, or using a data channel in the broadcast. Where possible we should try to use the timing information present in existing data channels, or potentially add it where missing.

Sean introduced the concepts of “studio time” and “stream time”. Studio time is aligned to the audio as it is heard in the studio, and would be the time reference for events triggered live in the studio.

For the purpose of synchronisation in RadioDNS applications, the timestamp need not necessarily be wall-clock time. It should be a common time reference shared by the broadcaster and the receiving device. We refer to this as the “stream time”, to indicate that it is a time reference aligned to the audio media.

The stream time may vary between broadcasters and bearers. For example a Radio 1 IP stream may contain a different stream time reference than Radio 1 broadcast on DAB.

The stream time should be a monotonically increasing value, and does not have to correspond to wall-clock time.

The problem for broadcasters to solve is how to map between studio time and stream time.

The radio receiver should be able to treat the stream time as an opaque value, using it to trigger events on the radio (e.g., display of a RadioVIS slide), or passing it to a RadioTAG service to indicate the time of a tag. It is up to the implementer of the RadioVIS or RadioTAG (or other) service to correctly interpret the timestamps. To do so, the service would need to know the broadcast service parameters as well as the timestamp.

- Each broadcaster can use their own stream time reference. Don't need to force all broadcasters to use the same
- Receivers can treat the time reference as opaque
- Broadcasters must be able to map a stream time to a known position in the audio

At what point in the production chain is the stream time reference inserted? For DAB, the time is set not by the broadcaster, but in the ensemble.

For IP streams we need to address how a common time reference will be set where there are multiple streaming servers, which may be re-started at different times.

Robin pointed out that we need a solution that works without changing the broadcast infrastructure.

The question of terminology was raised. Should the time reference embedded in the media be referred to as “stream time” or, as it may differ between bearers, “bearer time”.

Audio fingerprinting

Audio fingerprinting can achieve accurate synchronisation, but it is computationally expensive, both for the receiving device and the broadcaster. Hard to do for live broadcast as the fingerprint processing introduces delay.

Internet-based clock synchronisation

NTP synchronises a local clock to a remote clock. However, it doesn't prevent jitter or local clock drift between re-synchronisations.

Time signalling in broadcast platforms

James presented the capabilities of existing broadcast platforms and internet streaming format regarding embedded timing information. The slides for this presentation are available here:

<http://goo.gl/3s9C3>

Synchronisation of events to HTML5 video

Chris presented a recent BBC R&D project, LIMO ("Lightweight Interactive Media Objects"), which is a system based on HTML5 audio and video for delivering arbitrary timed metadata events and synchronising the triggering of these events to audio or video playback in a Web browser. The slides for this presentation are available here:

<http://radiotag.prototyping.bbc.co.uk/docs/LIMO-2012-08-24.pdf>

Review of RadioDNS draft document "Synchronisation of RadioDNS Applications to Audio"

The group expressed its thanks to Ben Poor for sharing this document. Reading through the document, the group made some observations:

- It is not clear whether this is a discussion document or the start of a specification. A discussion document that is separate from the specification could be useful at this stage - for example, to keep separate discussion of alternative approaches such as signalling delay in RadioEPG and device prediction of transmission delay.
- The introduction may be clearer if it introduced the various time references, e.g., studio time, stream time rather than causes of de-synchronisation
- Section 3.1 should be expanded to cover all bearers referenced in the RadioDNS Technical Specification (RDNS01). For each, we should specify how to obtain the stream time, and what level of resolution and accuracy we can expect to achieve
- In sections 3.2.1 and 3.2.2, should we use a RadioDNS-specific name instead of "Clock", such as "RDNS-Sync-Time".

Discussion and AOB

Also discussed were how we might use LIMO to update SVG widgets in a radio.