RadioDNS

— the hybridisation of Radio

Nick Piggott
Global Radio (UK)

A hybridised radio is a device that combines both Broadcast and IP technologies in a single functional unit, and where the broadcaster can instruct the client device to use either or both capabilities to deliver a homogeneous media experience to the listener.

The use of RadioDNS to combine Broadcast and IP technologies in hybridised radios will allow radio broadcasters to retain the economically attractive model of mass-market broadcasting, but also to enhance it by deploying innovative services faster and with lower development costs.

The first era of radio's technical development could be characterised as slow and steady. AM transmissions still cover the majority of the world's population, more than 100 years after the technology was introduced ... and 50 years elapsed between the first demonstration of FM modulation and its widespread adoption by broadcasters and listeners. In that context, one might consider Eureka-147 DAB, first demonstrated in the early 1990s, as still in the early stages of its adoption.

In comparison, whilst 20 years elapsed between the creation of IP (Internet Protocol) as a Transport Layer protocol in 1974 and the arrival of HTTP (Hyper Text Transfer Protocol) as an Application Layer protocol in the early 1990s, the number of Webservers (offering HTTP as an Application) grew from zero to 207,316,960 *(see Fig. 1)* over the period from November 1995 to February 2010 – in just over 14 years. *(Source: Web Server Survey, Netcraft Limited:* **http://news.netcraft.com**).

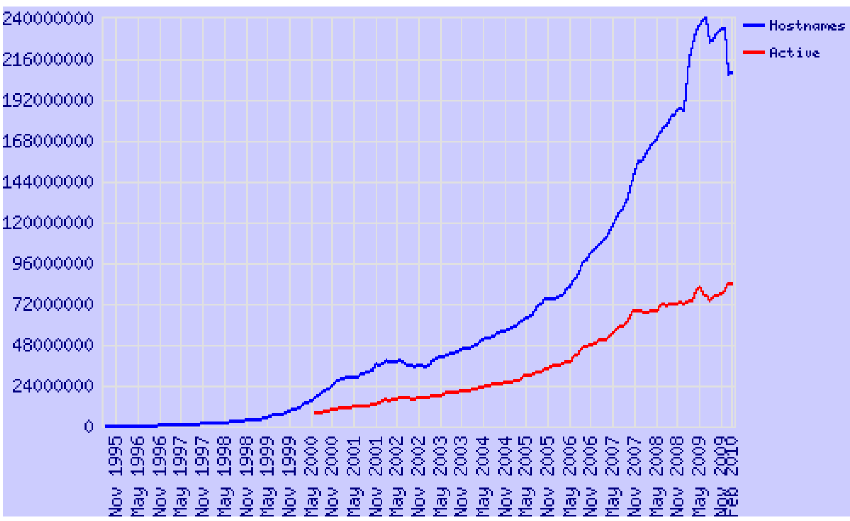

Figure 1
Growth of HTTP Webservers *(Hostnames = uniquely named sites. Active = number of active servers) © Netcraft Limited*

The Internet changes quickly. Radio broadcasters need to understand how that speed of change affects the medium, and how to respond to it effectively and without undermining the constancy and consistency that characterise radio and that are cherished by many of its listeners.

The commercial radio environment

The pressure for change is particularly acute in the commercial radio sector, which makes up roughly half the market (by listening hours) across Europe. Whilst public-service broadcasters need to demonstrate that they are continuing to deliver value by delivering solid audiences, commercial broadcasters must do more than that.

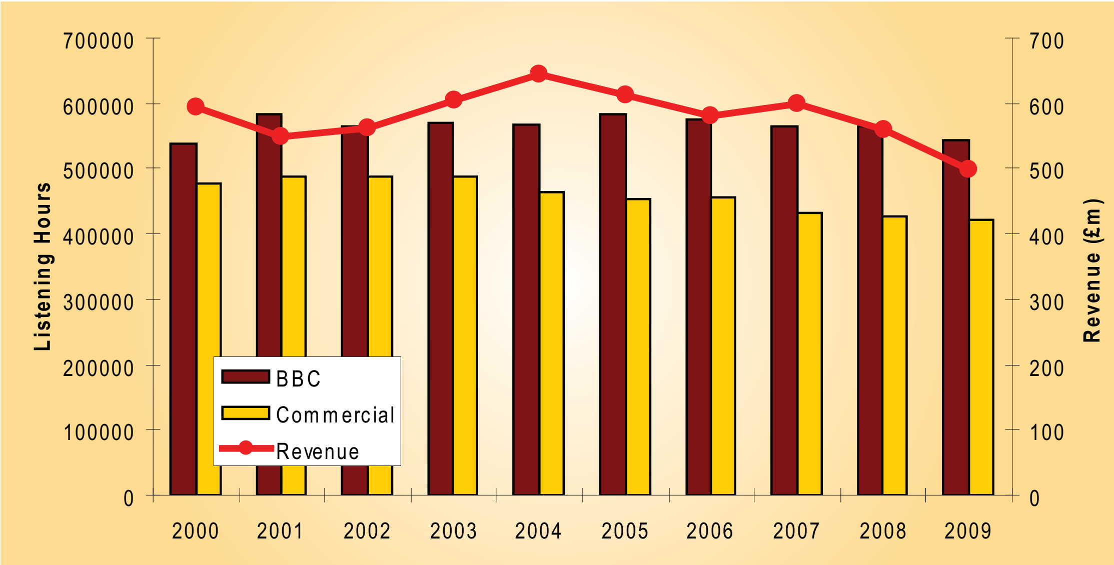

Figure 2
UK listening hours and commercial radio revenue

Using the UK market between 1999 and 2009 as an example *(Fig. 2)*, the overall listening hours have declined by just 1.9% – from 1,006,330m in 1999-Q4 to 987,584 in 2009-Q4 *(Source: RAJAR).* However, commercial radio has recorded a larger decline of 10.4%, dropping from 469,721m in 1999-Q4 to 421,063 in 2009-Q4 *(Source: RAJAR).* Worse still, commercial radio revenues have fallen from £594m to £497m (–16.4%) over the same period, and are down (–23%) from a peak of £645m in 2005 *(Source: Radio Advertising Bureau, UK)*.

All “traditional” media (newspapers, radio, television, cinema) have seen revenues decline as advertisers have increased the proportion of their overall marketing budgets spent on Internet-enabled marketing. That might be display advertising on websites, direct marketing using e-mail, dedicated brand websites and, most recently, investing in social networking sites. The spending on online advertising has exceeded that on radio advertising in the UK since 2005 *(Source: Internet Advertising Bureau, UK)*, although many radio broadcasters also have successful online presences which attract online spend.

This changing commercial environment is a significant factor in affecting how commercial radio broadcasters are considering their strategic directions and asking how technology can help both to retain audiences and protect revenues.

Digital Radio

The speed of change, and the commercial pressures that accompany it, explain why Digital Radio can mean different things to different people.

To start with, there are a number of different broadcast standards which can carry audio / radio services:

- DAB (Eureka-147 with MPEG2-II audio encoding);
- DAB+ (Eureka-147 with HE AAC v2 audio encoding);

- DRM (Digital Radio Mondiale);
- IBOC (of which iBquity's "HD Radio" is the most prominent technology);
- DVB (all variants : -T, -T2, -H, -C and -S).

But there are also people who consider that the digitization of radio can take other forms:

- Radio stations having websites;
- Presence on other websites, such as social networks;
- Streaming audio, delivered over fixed/wireless IP networks to personal computers;
- Streaming audio, delivered over UMTS networks to mobile phones;
- Streaming audio, delivered over wireless IP networks to dedicated appliances ("WiFi Radios");
- Downloaded audio, delivered to personal media devices ("Podcasting");
- Presence on other IP-connected devices such as games consoles.

These broader definitions of "the digitization of radio" encompass a more commercial understanding of radio as an advertising-funded medium, or a broader concept of how public-service broadcasters can attract audiences. These views consider the audience as a whole, regardless of whether they listen to the existing time-linear radio service, or devote media time to the radio station in some other way.

These various interpretations, and in some cases choices of technology, can have a significant affect on the cost of distributing a radio service without necessarily having a clear benefit in terms of listening or revenue. A broadcaster making a cost-benefit analysis of the options open to them has to consider over what timescales they need to see the benefits, the scale of the benefits, and to what extent they might or might not materialise.

The benefit of the Internet

The Internet, or rather the open connectivity that it provides, has stimulated innovation and development of new applications, services and devices at a tremendous rate. The lowering of traditional

Abbreviations

AAC	Advanced Audio Coding
DAB	Digital Audio Broadcasting (Eureka-147) http://www.worlddab.org/
DAB+	DAB using the AAC codec
DNS	Domain Name System
DRM	Digital Radio Mondiale http://www.drm.org/
DVB	Digital Video Broadcasting http://www.dvb.org/
ECC	Extended Country Code
EId	Ensemble identifier
EPG	Electronic Programme Guide
FCC	Federal Communications Commission (USA) http://www.fcc.gov/
FM	Frequency Modulation
FIG	(DAB) Fast Information Group
FQDN	Fully-Qualified Domain Name
GSM	Global System for Mobile communications
HE AAC	High Efficiency AAC
HTTP	HyperText Transfer Protocol
IBOC	In-Band On-Channel
IANA	Internet Assigned Numbers System
IEEE	Institute of Electrical and Electronics Engineers (USA) http://www.ieee.org
IETF	Internet Engineering Task Force http://www.ietf.org/
IP	Internet Protocol
ISP	Internet Service Provider
NIC	Network Information Center
PI	Programme Identifier
QoS	Quality of Service
RT+	RadioText Plus
RDS	Radio Data System http://www.rds.org.uk/
RFC	Request For Comments (IETF standard)
SCId	Service Component identifier
SId	Service identifier
UMTS	Universal Mobile Telecommunication System http://www.umts-forum.org/
URI	Uniform Resource Identifier
URL	Uniform Resource Locator

barriers to entry, and the removal of most of the regulation that hinders change, has accelerated change and created a competitive environment that sees success and revenues change dramatically and quickly.

The Internet consists of two distinct elements, with distinct funding for each.

- The ***infrastructure*** – providing physical connectivity – is largely funded by the connected parties. Businesses (including web-enabled services) pay for their industrial-grade connectivity to provide services into the network, while consumers pay a provider for connectivity to receive the services.
- The ***services*** – providing content and a *reason* for consumers to connect to the network – are almost wholly funded through advertising, either using a display advertising model or direct marketing model, or a combination of the two. Successful models where the consumer pays for a service are rare.

It is of course worth mentioning that much speculative development of new services is funded through venture capital on the assumption that: (i) those services that are successful have returns far in excess of the investment needed to start them up, (ii) the entry costs and barriers are low, and (iii) the cost of terminating a project is virtually nil.

The Internet model is almost diametrically opposite to the traditional broadcast model *(see Table 1)*. It becomes clear that there are very attractive aspects to the Internet-connected model. It's possible for an established broadcaster to divert funding and resources to Internet-enabled activities and see the benefits very quickly, and to do so without significant longterm commitment or risk. The speed of development of new products and services is limited by the end browser / computer, which itself evolves extremely quickly. This allows broadcasters to more rapidly address competitive threats. The commercial models for online services are well understood, and can be sold either in combina-

Table 1
Broadcast vs. Internet models

	Broadcast	Internet
Infrastructure funding	Broadcaster (Public Funding / Advertising)	Service Providers (Advertising), Network operators (Connection Fees)
Barriers to entry	Spectrum availability, Regulatory requirements and significant funding needed	Funding
Implications of exit	Outstanding infrastructure liabilities, building leases and Broadcast licence fees	Few
Speed of development limited by	Consensus, regulatory speed, technical standards and end-device turnover	Skill, end-device turnover (PCs, browsers) and connectivity speed
Audience addressability	Geographic and Demographic (through the proxy of the radio station's output appealing to particular segments of the population)	Individual [a]
Measurability	Diary-based surveys of samples of listeners and metered measurements of samples of listeners	Individual [a]

a. Individual is assumed to be "individual as a person", which requires the individual to identify themselves across multiple devices; the concept of a "log-in" or "authentication". In practice, much individualisation on the current internet addresses the "individual user of a device", through the use of cookies and similar techniques. Where a person uses multiple devices, the network and service provider consider them to be different "unique users".

tion with traditional radio advertising or entirely separately, providing some isolation from the vagaries of having a single revenue stream.

Finally, of course, the broadcaster has readily available resources to market their online services at zero cash cost, simply by promoting to the radio station's existing listeners. Some radio stations are able to convert 25% or more of their weekly listeners into users of their websites, although this conversion rate varies considerably according to the target demographics and geography of the radio station.

There are, of course, drawbacks to the Internet. As a technology, it's not well suited to broadcasting, as it was inherently designed to transport packets between two points, not distribute packets from one point to many. Techniques like multicasting (IETF RFC 3170 / 3171) are applications deployed onto router infrastructure to mimic broadcast-like behaviour, but require a consistent deployment within a network and co-ordination between networks to make them work as required. Similarly, wireless technologies such as WiMax (IEEE 802.16) allow for high data rates to mobile terminals, but only in specific and relatively low noise environments.

The density of a Broadcast network transmission system is generally dictated by the planning criteria applied to it, in particular the outgoing interference to other spectrum users. In comparison, the density of a wireless IP network is imposed by the restrictions in power that can be generated on the uplink channel by the device, which is often battery-powered. This leads to infrastructure (base station) requirements far in excess of those required for traditional Broadcast to cover the same geographical area. Simon Mason of Arqiva (UK) presented – to the 2008 WorldDMB General Assembly – an indicative figure of 117 WiMax base stations to cover a 5km radius area of Central London (UK).

Whilst the discussion about whether the use of IP technology often centres around current and predicted technical costs, this may be missing the significant ***economic*** difference between IP distribution and wireless broadcast distribution. IP networks almost certainly place one or more intermediaries (ISPs – Internet Service Providers) between the broadcaster and listeners, who control the infrastructure over which data flows. The listener has to pay a subscription to the ISP, which is not the universal free-to-air model of radio broadcast supported by governments in most countries.

In the UK, broadcasters have reduced the cost of distribution via the Internet to virtually nothing by peering directly into the main ISPs, thus eliminating "transit bandwidth" (for UK listening) almost entirely. However, this economic benefit is only available by joining LINX (The London Internet Exchange – www.linx.net) and only on agreement with the ISPs on a bilateral basis. In addition, the addressing space for multicast applications (in IPv4) is relatively restricted, which requires ISPs to be selective about which content services they enable for multicast within their network.

This is where the real economic benefit of Broadcast, and the risk of IP, lies. The Broadcast model, as envisaged by most governments, grants spectrum to broadcasters to provide a universal, ubiquitous, free-to-air service, with high quality of service, directly to the listener. To achieve some of the same objectives using Internet distribution requires numerous bilateral commercial agreements with ISPs, securing bandwidth, quality-of-service and multicast address assignments. There is indication that, whilst basic "best efforts" unicast distribution may continue to be freely provided, content distribution businesses requiring a specific level of service will have to pay a premium for it, with no guarantee that it will be available, particularly where a resource such as multicast address space, is limited.

Figure 3
DAB Radio – digital technology in a traditional form
Photo: Author's own

If a broadcaster can reduce its reliance on consumption of the economically "expensive" activities on IP networks (QoS, multicasting), he can reduce the cost of distribution, and can use the network to carry only traffic which is delivering more value than the broadcast content. This might be complementary or on-demand content, content that appeals to smaller numbers of listeners ("niche" services), or transactional information providing the broadcaster with information on a listener's preferences, behaviour or desire to purchase.

Therefore, rather than having two competing technologies, we appear to have two technologies that are optimized for complementary tasks and have different economic models:

- **Broadcast –** is suited to the mass distribution of (digital) content to large numbers of low-complexity / low-cost receivers, on an entirely free-to-air model. The technology is stable and devices have long life cycles.
- **IP** – is suited to the targeted distribution of content or transactional information to a fewer number of more complex devices, connected to subscription-funded networks. The technology changes rapidly, and devices can be obsolete within 3-5 years.

With these circumstances in mind, how can the concept of a radio that combines both broadcast and IP technologies be realised?

Hybridised Radio

A hybridised radio is a device that combines both Broadcast and IP technologies in a single functional unit, and where the broadcaster can instruct the client device to use either or both capabilities to deliver a homogeneous media experience to the listener.

Arguably, hybrid devices have been in existence for some time. Nokia have released GSM mobile phones that include FM radio functionality for some years, but there has been no link between the two functions. The IP-connected device had no awareness or understanding of the FM radio function, and the FM radio had no visibility or understanding of the IP connectivity it had available to it.

Two projects have come close to Hybridised Radio, but failed in significant technical areas.

Nokia Visual Radio was a project instigated by Nokia, which launched in 2006. The FM Tuner application on the handset was enhanced to include the Visual Radio functionality, which in turn allowed the broadcaster to deliver synchronized visual content via IP to accompany the FM radio broadcast. The business model initially envisaged networks charging subscribers a modest additional fee to access the service, although this ultimately failed.

Whilst it was a combination of broadcast content plus IP-delivered content, it failed to be a true hybridisation as it required Nokia to manually maintain a table of station-specific identifiers (VRIDs – Visual Radio Ids) and for the listener to use a specific Station Directory on the device in order to tell the client software where to locate the visual feed; the broadcaster could not signal it directly, and the client could not determine it by examining the broadcast signal.

RT+ (RadioText Plus) is an extension to the existent FM-RDS RadioText application (specified in IEC 62106 Ed.2), that was jointly proposed by Germany's *Institut für Rundfunktechnik* (IRT) and Nokia. RT+ allows the broadcaster to transmit additional information in RDS Data Groups that can indicate to the receiver parts of the RT text

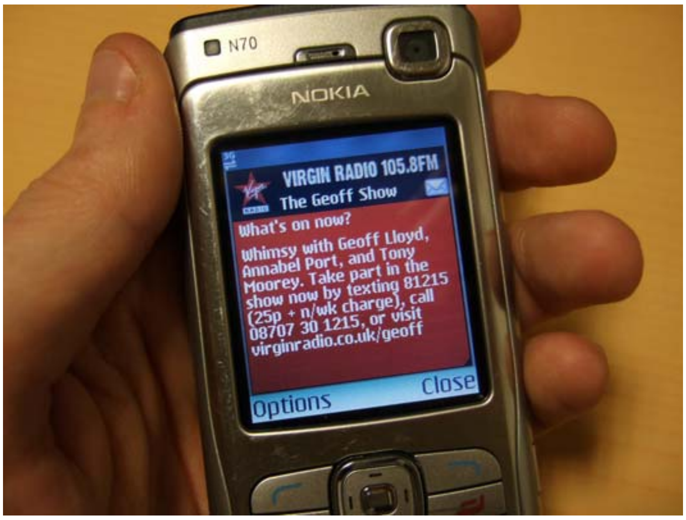

Figure 4
Virgin Radio 105.8FM on Nokia Visual Radio.
Photo: James Cridland under Creative Commons licence 2.0 Generic

string to be interpreted as machine-readable data (“metadata”). For instance, the human readable text string of “Now Playing: Starship with Nothing’s Going To Stop Us Now” could be marked up using RT+, such that the client would know that the artist was “Starship” and the title of the song was “Nothing’s Going To Stop Us Now”. It was intended that the client could then offer the user various options, such as purchasing the song. This was also an example of using IP connectivity to enhance broadcast radio, but failed to be truly hybridised because the client passed all information back to a single service provider, rather than locating the originating broadcaster’s presence on the Internet, and passing the information to them.

The advent of “WiFi” radios, combining both broadcast radio receivers (FM / DAB / HD Radio) and 802.11 wireless IP technologies should have been an opportunity to create a hybrid radio but, in the majority of cases, the technologies have been implemented in discrete functional modules, and the only shared resources are the screen, user interface, power supply and audio output stage(s). Most radios combining functionality in this way still require the user to make a high-level decision to choose a *platform* (FM, DAB, HD Radio, Internet) and then navigate a station listing which is specific to that platform. This can lead to confusion where one radio station is available on more than one platform, which is compounded when the “favourite stations” function also maintains separate lists on a per-platform basis.

Connecting Broadcast and IP

The principle of hybridisation is that the end-user should be unaware of how broadcast and IP are being combined to create a homogeneous experience, and it should happen automatically. To achieve this outcome, two steps are required:

Step 1: Identify which radio station the user is listening to;

Step 2: Locate the IP-delivered services for that station.

It’s conceivable to combine these two steps into one composite method. As an example, it might be possible to transmit, as part of the broadcast channel, the URI (Uniform Resource Identifier; IETF RFC 3305) or FQDN (Fully-Qualified Domain Name; IETF RFC1035) of the broadcaster. That could be accomplished using methods such as ODAs (Other Data Applications) in FM-RDS or FIG (Fast Information Group) signalling in DAB.

However, proposing changes to existing standards, and implementing those changes in both the transmission chains and receivers is time-consuming and expensive. As noted earlier, radio needs to use IP to deliver enhanced services to market faster, and this suggests using an alternative approach to solve Step 1, rather than time-consuming standardization work.

Abstracting the problem, the first issue requires the *unique* identification of the originating radio service. This requirement can be met, in most cases, through metadata already present in many existing broadcasts:

- **FM-RDS:** ECC country code and PI programme identifier;
- **DAB:** ECC country code, EId (Ensemble identifier), SId (Service Identifier) and SCIds (Service Component within the service);
- **HD Radio:** FCC Facility Code;
- **DRM**: SId (Service Identifier).

These codes ought to uniquely describe a service, and thus appear to meet the requirement of Step 1.

To achieve Step 2, a centrally-managed web-service at a prescribed URL could return the actual location of a broadcaster’s IP services in response to being provided with the identity of the radio broadcaster. However, the service would need to be scalable, and have high reliability and uptime, taking into account that every connected radio in the world would access it every time it acquired a new radio service. There would also be issues of network latency and traffic, unless the database was distributed across different physical locations.

However, this problem is almost identical to that faced by every connected device when trying to translate "human-readable" URLs to the underlying IP addresses. To handle the demands on the network that this would create, DNS (Domain Name System; IETF RFC 1034 & 1035) was designed with scalability, distribution and security in mind. Each connected device contacts its preferred DNS server using a protocol optimized for name resolution. In turn, each DNS server contacts its parent and peers to synchronize changes to the DNS database, with the hierarchy topped by a number of "root" servers co-ordinated by IANA (Internet Assigned Numbers Authority – **http://www.iana.org**). DNS servers typically hold millions of entries, often on a cached basis, using hinting information provided in the domain records themselves. The root servers can reference the entire universe of domain names, which are themselves managed by NIC (Network Information Center) operators who maintain root servers for specific domains such as .com, .org, .co.uk etc.

In addition, the DNS allows for relatively rich sets of information to be held on each domain record. A non-exhaustive list of record types for a domain is shown in the Table below:

Record Type	Function
A	Numeric IP address (IPv4)
MX	Server handling mail for this domain
CNAME	Canonical Name
SRV	General Service Location record

Therefore DNS seems to offer the scalability, reliability and distributed nature that achieving Step 2 requires, with the added benefit that DNS client software stacks are readily and widely available, and universally implemented in all connected devices. The record types allow for a wide variety of services to be located.

However, DNS only supports the lookup of FQDNs. This poses the question – *is it possible to map the station identification data acquired in Step 1 into a format that looks like an FQDN?*

An obvious solution, and one used in parts of various DAB-related specifications, is to separate the individual elements using "." characters. This would lead to the following mapping process:

RDS-PI code: C586

RDS-ECC code: CE1

Constructed FQDN: c586.ce1

A similar example for DAB might be:

DAB-SCIds: 0

DAB-SId: c586

DAB-EId: c185

DAB-ECC: ce1

Constructed FQDN: 0.c586.c185.ce1

This begins to look like a partial solution. However, the root element of this translation methodology would be either RDS-ECC or DAB-ECC, of which there are potentially many variants. This would require the top level registration of every conceivable permutation of ECC in order to protect and delegate the namespace.

One way to achieve this protected namespace, and to reduce the registration burden, would be to append a standard domain to the FQDN constructed by concatenating the service identifiers, such as radiodns.org.

Now our examples become:

Constructed FQDN: c586.ce1.radiodns.org

Constructed FQDN: 0.c586.c185.ce1.radiodns.org

These are now valid Fully Qualified Domain Names delegated in the zone radiodns.org. The root name server for radiodns.org can hold all the required variations of service identifiers.

Therefore it appears that there is a standard methodology that can be applied to FM-RDS and to DAB to translate service identifying parameters into a valid FQDN, and that the same methodology can be adapted for similar broadcast systems. This FQDN can exist as a domain record, within the radiodns.org zone.

The client need only apply the method to the service identifiers to create an FQDN representing the currently received service, pass this FQDN to DNS using its existing DNS client software, and receive an A record containing an IP address in return. It has now completed Steps 1 and 2, and has a valid IP address with which it can connect to the broadcaster via IP. This has all happened automatically, within a few seconds of tuning into the broadcast service.

Prototype implementation

This concept was prototyped in early 2008 by the Creative Technology Team at Global Radio, in collaboration with the BBC's R&D team. In the course of prototyping, a number of weaknesses in the theoretical architecture were identified and addressed.

Firstly, it was identified by examination of the records published by OFCOM (UK Regulator – **http://www.ofcom.org.uk**), concerning the PI codes in use across the United Kingdom, that the PI codes were not unique. It's not clear if this was by design or by error. In order to widen the available namespace and reduce the chance of collisions, it was decided to prepend the FM frequency (as a five digit integer) to the constructed FQDN.

Secondly, it became clear that transmission of ECC codes was inconsistent, and that timely acquisition of the code was reliant on broadcasters transmitting it relatively frequently. However, it appeared that many devices might already know which country they were located in by other means, including through user configuration or inspection of information from a GSM base station, so an alternate FQDN construction method was provided for FM that allowed the country to be provided as a two character ISO code in the absence of an ECC code.

Thirdly, it was decided that the semantic structure of the FQDN would be improved by making the platform a high-level identifier in the FQDN, placed directly after the radiodns.org element. This removed any confusion over whether an ECC code is part of an FM or a DAB service, as the ECC value is common to both platforms. It also allowed for a logical delegation or partitioning of the zone in the future if required.

Finally, minor amendments were made to the methodology for constructing FQDNs for DAB services to encompass any arbitrary service, not just audio services, including data services delivered by X-PAD (although this is a semantic issue that does not immediately appear to have any real-world application).

The issue of application discovery was addressed, with reference to the state-of-the-art methods used in the IP domain.

The simple theoretical method holds a single A record, and returns a single IP address. However, this still requires the client to negotiate capabilities with the broadcaster's servers, which could become costly if there are a broad set of applications available. It also doesn't allow the broadcaster to delegate functionality to third parties, as all traffic is directed to them in the first instance. Additionally, IP addresses can be quite dynamic, and whilst broadcasters can often directly control their own domain records, there would be an administrative overhead and timing issue making changes to IP addresses in records in the radiodns.org zone.

For these reasons, the concept was amended. Instead of returning A records, the domain record in the radiodns.org zone would hold one single CNAME record. A CNAME is a form of redirection, indicating to the client that they should reform their request using the FQDN returned in the CNAME response, rather than the originally requested FQDN. This dramatically simplifies the maintenance of the radiodns.org zone, by essentially reducing it to a simple re-direction service, as illustrated in the Table below.

Step	Data
Acquire Station Identifiers	PI=C586, ECC=CE1, f=95.8MHz
Construct FQDN	09580.c586.ce1.fm.radiodns.org
Pass FQDN to DNS server	CNAME = rdns.musicradio.com

Once this initial re-direction has been performed, all further DNS queries are against the broadcaster's own domain records, and radiodns.org plays no further part.

DNS can also help simplify the process of determining which applications the broadcaster is providing over IP. The client needs to determine which of the applications it can handle are supported by the broadcaster, and how to contact the specific server(s) handling that application. This problem is addressed in DNS through use of SRV records (IETF RFC 2782), which act as pointers to specific services.

The client can query DNS for a specific SRV record. If the record is absent, the broadcaster does not support it and the functionality can be disabled for the user. If the record is present, it tells the client device which server and port to contact. This has the added benefit of allowing different applications to be handled not just by different servers, but potentially by entirely different organizations, allowing for complete delegation of application support to third parties.

The SRV records use a maintained list of applications, and new applications can be added on request to the list maintainer (**http://www.dns-sd.org/ServiceTypes.html**).

As part of the prototype service, three applications were tentatively defined, and registered as follows:

Application Name	Purpose	SRV identifier
RadioVIS	Visual enhancement (text and pictures)	_radiovis._tcp
RadioEPG	Electronic Service and Programme Guide	_radioepg._tcp
RadioTAG	Simple Interaction Framework	_radiotag._tcp

The specifications of these applications can be found at **http://radiodns.org/docs**. Further application types can be freely defined, and deployed on a standardized, proprietary or ad-hoc basis, as the reference function of DNS is application-agnostic.

This SRV record discovery phase extends the prior example as follows:

Step	Result
Acquire Station Identifiers	PI=C586, ECC=CE1, f=95.8MHz
Construct FQDN	09580.c586.ce1.fm.radiodns.org
Pass FQDN to DNS server	CNAME = rdns.musicradio.com
Construct FQDN for RadioVIS SRV query	_radiovis._tcp.rdns.musicradio.com
Pass SRV query to DNS	SRV = 80 vis.musicradio.com (port / server)
Pass FQDN returned in SRV record to DNS	81.20.48.118
Connect to RadioVIS server	81.20.48.118:80

The client now has a connection to the required server, and continues as specified in the RadioVIS specification.

The prototype system was hosted on two virtual servers (ns1.radiodns.org and ns2.radiodns.org) using the Elastic Computing Cloud service provided by Amazon Web Services. The implementation used a standard installation of a BIND DNS server. Maintenance of the radiodns.org.zone file was by hand.

During the life of the prototype, broadcasters in more than 10 different countries, across three continents, asked to trial the service and were given test entries in the system. At the beginning of 2010, over 500 radio stations were represented in the database.

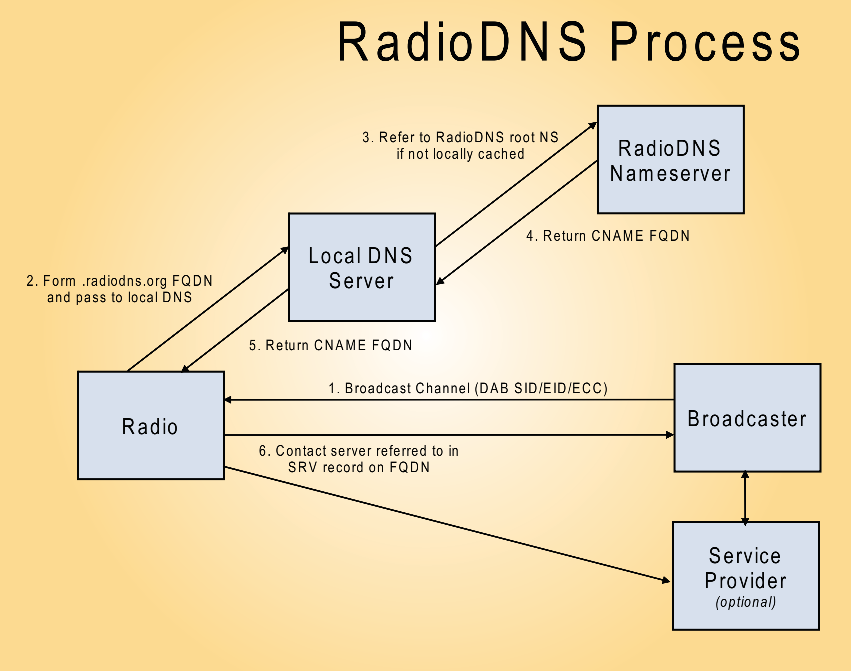

Figure 5
RadioDNS Process in functional blocks

Published specifications

The proposed RadioDNS concept and specification was first presented to the WorldDMB Technical Committee meeting in May 2008. The draft documents for the RadioDNS specification and its associated applications (RadioVIS, RadioEPG and RadioTAG) were published on the project website shortly afterwards. At the time of writing, the status of these specifications is:

Specification	Version	Status
RadioDNS	0.61	Stable, in production use
RadioVIS	1.00	Stable, in production use
RadioEPG	0.61	Draft
RadioTAG	0.61	Draft

Other applications can be created, either as part of the RadioDNS project or outside it.

Production implementations

The most prominent implementation of RadioDNS to date has been in the Sensia Connected Radio *(Fig. 6)* from PURE (a division of Imagination Technologies, UK). This innovative radio combines FM, DAB and IP connectivity (using 802.11 wireless technologies).

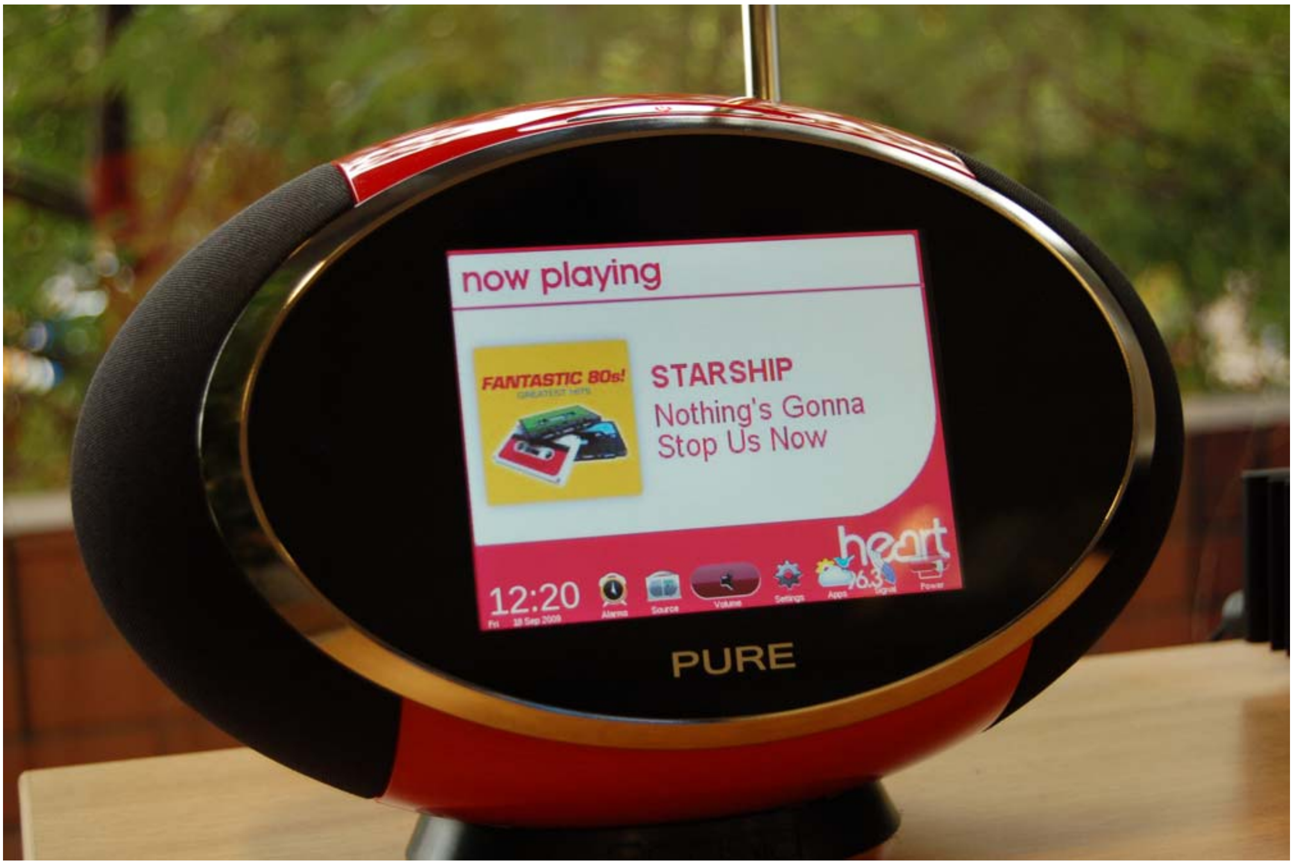

Figure 6
PURE Sensia *by Byrion Smith under a Creative Commons 2.0 Generic licence*

The Sensia is an example of a hybridised radio which implements the RadioVIS application, when receiving DAB radio services. Other manufacturers may bring radios to the market implementing other applications.

The Sensia uses the RadioDNS process to connect to a RadioVIS server to receive images and text. A number of broadcasters have production implementations of RadioVIS servers, including Global Radio and others in various countries. In addition to providing visuals for radio devices using RadioDNS, the Global Radio service is also used to enhance mobile phone applications and browser-delivered streaming services with visual information, thus extending the value of the investment in content production.

Formalisation

The RadioDNS prototype demonstrated that the technology works in practice, and it's a testimony to the thoroughness of that prototype that a manufacturer (PURE Digital) was sufficiently confident to launch a commercial product using it. The widespread interest from broadcasters and manufacturers secured the decision to turn the prototype into a production service.

With that in mind, RadioDNS was formed on 16 February 2010 as a not-for-profit organization, with the objectives to

Nick Piggott is Head of Creative Technology at Global Radio (UK). He has spent the majority of his career working in the commercial radio industry, initially in programming and then moving to digital radio and new technology in the late 1990s. In 1999 he built up a small multi-skilled team to launch Planet Rock and Core as the world's first commercial DAB Digital Radio stations. Under his leadership, that team has consistently delivered innovative new services to radio stations, including groundbreaking interactive text message services, music downloading, mobile and multi-platform applications, and visualised radio services. The team has won numerous awards for innovation over the last decade.

In 2008, Mr Piggott was one of the founders of the RadioDNS project to converge broadcast radio and IP. He is an active member of the WorldDMB forum, and talks globally about the opportunities for radio in the digital age.

promote RadioDNS and associated technologies and applications, and to provide the RadioDNS service on a permanent basis. It is expected that the RadioDNS Steering Board will issue an RFP in due course for a service provider to operate the radiodns.org root servers and manage the addition of new radio stations to the service.

RadioDNS remains an open and collaborative project, which welcomes new supporters and members from any part of the radio business. The project website is at **http://radiodns.org**.

Conclusions

The use of RadioDNS to combine Broadcast and IP technologies in hybridised radios will allow radio broadcasters to retain the economically attractive model of mass-market broadcasting, but enhance it by deploying innovative services faster and with lower development costs. These new services could be of significant commercial value to commercial broadcasters.

This version: 16 March 2010

Published by the European Broadcasting Union, Geneva, Switzerland

ISSN: 1609-1469

Editeur Responsable: Lieven Vermaele

Editor: Mike Meyer

E-mail: **tech@ebu.ch**

The responsibility for views expressed in this article rests solely with the author