

Linux at the BBC

What do *Strictly Come Dancing*, the shipping forecast and open source have in common? They're all big at the BBC, as **Alex Singleton** finds out.

The British Broadcasting Corporation is the biggest broadcaster in the world. It spends more than £4bn a year and employs 28,000 people. It transmits eight UK television channels, six overseas channels, countless radio stations in over 200 countries, and has a website with over two million pages. As you'd expect, the BBC's technical requirements are rather different from those of the average business IT user. Its support contract with Siemens is worth nearly £200m a year, for example, and a technical fault can mean millions of people looking at a blank screen, either on their TV or their computer.

The BBC's research and development team is based at Kingswood Warren, a large country mansion in Surrey. For the staff at Kingswood, Linux is the operating system of choice. They prefer it as a development platform and they find that demanding software runs better on it than it does on Windows. Here Linux-based applications are being developed that push the operating system's use out of its traditional niches and into the heart of the BBC's work.

Better for broadcasting

The BBC's empire stretches computer equipment to its limits: if a studio is being used to record high-definition television coming in from multiple cameras, it needs systems that can cope with four or five 1.5gbps streams of data simultaneously – and the Kingswood team find Windows increasingly unable to cope – but with Linux, the BBC is able to take advantage of the system's modular design and choose parts that won't buckle under the strain.

"Windows uses Microsoft's *NTFS* filesystem," says BBC senior software engineer Stuart Cunningham, "but it isn't very good. It can't cope with all the data we're trying to simultaneously save. In theory, you can use other filesystems under Windows, but it's difficult and causes lots of problems. Under Linux, we tried all the different filesystems, found that *XFS* was the best for our needs, and just slotted it into place."



› The BBC's David Kirby believes Linux-based technologies can deliver significant cost savings.

Cunningham works on the department's *Ingex* project, which is a Linux-based storage system that cuts days off the editing process. Ordinarily, broadcasters record on to tape. The problem is that a day's



filming with several cameras can easily result in 35 hours of recorded material. If recording takes a week, programme-makers can wait days for the tapes to be inputted into an editing system. *Ingex* eliminates that process by feeding the camera output directly into a Linux-based computer with a gigantic hard disk. It's currently at a prototype stage but it has been used for real on broadcast programmes, including *EastEnders* and *Dragon's Den*.

Hackability

"Linux gives us full control," says Cunningham's colleague David Kirby. "As it's open, we can tune and enhance the system so we get all the performance we need. One of the design considerations was that different programme-makers use different editing suites." Some users prefer *Avid* while others use *Final Cut Pro*. *Ingex* stores video using a combination of the open standard MXF wrapper, which contains Advanced Media Format data. The problem with this is that *Final Cut Pro* doesn't understand files stored in an MXF wrapper. Rather than store files in two formats, with all the extra processing time and hard disk space that that would entail, the Kingswood team have created a Linux-based virtual filesystem. When *Final Cut Pro* users try to access the files, the virtual filesystem kicks in and ditches the wrapper, simply showing the Advanced Media Format file inside.

The *Ingex* team benefited hugely from the open source community: some 90% of the code used by the system was pre-existing open source code. They were also able to pay external developers to work on an open source version of one of the compressors they used, so they wouldn't have to keep paying for a proprietary version. For a high-definition programme with four cameras, *Ingex* uses two dual-processor Core Duo PCs. Commodity equipment is



› Kingswood Warren, a Gothic mansion completed in 1837 and site of major BBC innovations.

› The set of *Dragon's Den*, where production staff have saved days thanks to *Ingex*.



essential to keep the price down – at the moment the BBC uses a mix of Dells and HPs. The PCs mainly run on OpenSUSE 10.3 but the team are testing the software on other Linux distributions, particularly Ubuntu and Red Hat too.

David Kirby is hoping that the final product will end up being run on a distribution that's free as in beer. Having to pay for the operating system goes against one of the main considerations of the project, which is to make the system as affordable as possible. Ultimately, though, that choice will depend on whoever ends up supporting the systems (the research and development team's remit doesn't involve spending its days in television studios maintaining finished products).

Cutting needless costs is a central reason why the BBC is looking to Linux and open source. As viewers increasingly expect to freely download BBC programmes over the internet, proprietary technologies, especially ones that demand per-user royalties, can damage the economic viability of BBC services. "Our business model isn't compatible with per-user royalties," says Tim Borer, the father of a new BBC technology called Dirac.

Codec kaleidoscope

"At the moment we have a lot of video material on the web and it's encoded in different formats," says Borer. "We've got Windows Media, RealMedia, Flash and a little QuickTime. All those types of content have to be managed, and we have to pay royalties on some of them. We have to pay for RealMedia, for example, based on how many people are streaming video simultaneously. It's not just that this takes money away from programme-making, it also means that we have to create complicated auditing systems to calculate precisely how much royalty money we should be paying. When we use RealMedia, we have to implement a 'back channel' from the player back to the server, and that means we have to have a mechanism for getting the back channel to work properly – for example, through firewalls. It's a lot of hassle and is completely unnecessary."

So Borer set up the Dirac project to cut out the waste. "Dirac has been a research project, and it's now approaching a 1.0 release," he explains. "It's a free and open source codec, and what's key is that it's patent-free. That's important to us, because there's absolutely no point in developing another patented codec with royalties. After all, there are perfectly good codecs already if you don't mind the restraints that occur because of the royalties."

Dirac has another advantage: "At the moment," Borer continues, "we're forced to use a whole variety of video formats, with content often in multiple formats so that BBC content works on everyone's computer. This means that BBC staff are wasting time simply converting files. We've used RealMedia for a long time because it's cross-platform, but a lot of end users don't like

RealMedia because they find it intrusive and refuse to install it. So we have to support Windows Media Player, too. One of the things we're trying to do with Dirac is to provide a cross-platform format that everyone can use and we can escape the royalties. We'd like to see all of the BBC's video material encoded in Dirac and encoded only once."

Value for money

Linux users have often favoured the royalty-free Ogg Theora format, but Borer thinks it isn't good enough. "Most public service broadcasters around the world are in a similar position to us in that we don't want to pay patent royalties, but until Dirac there haven't been any real alternatives. Theora has been around for a while but it's based on very old technology, and it's not very good at compressing video. Yes, it's better than MPEG2 (the compression used on Freeview and on DVDs), but time has moved on. It is not as good as H.264 though, which is being used by high-definition broadcasting."

The issue of high-definition broadcasting is critical. As computer screens become ever larger and broadband speeds increase, people are starting to expect TV to be available on their PC, and not just with pixels-in-sight resolution. But for high-definition internet broadcasting to be practical over the sort of broadband connections that people are likely to have in the foreseeable future, it's essential to get the compression right.

"Dirac will compete with H.264 on a technical level," says Borer. "We've designed it to be very simple. It's much easier for companies to implement

DAVID KIRBY ON OPENNESS

"As Linux is open, we can tune and enhance the system to get all the performance we need."



› Brandon Butterworth led the development team that created a Flash version of the BBC's *iPlayer*.

Platform-neutral content

"The original version of *iPlayer* used Windows Media Format and digital rights management," says the BBC's Brandon Butterworth, "and there was a big fuss." He's right: it led to demonstrations outside the BBC's Television Centre in London and over 16,000 people signing a petition in protest. After constructive dialogue between the Open Source Consortium (a UK-based trade association) and the BBC Trust, which governs the broadcaster, the BBC Trust listened to the criticism and launched an investigation into how the state of affairs could be improved.

It's not common knowledge, but the BBC has an obligation to encourage open formats. Associated with its Royal Charter is a 2006 agreement with the UK government that stipulates that the broadcaster "must pay particular attention to the desirability of supporting actively in national and international forums the development of 'open standards' (that is to say, technologies where opportunities to participate in their creation and

to use them are made widely available, free of charge or on terms that must be fair, reasonable and non-discriminatory)."

So the BBC Trust prodded BBC management, who in turn looked to the R&D team at Kingswood Warren for help. Brandon Butterworth explains: "People thought it would be really difficult to run on all platforms, that it would take years to develop. But I led a team to do a Flash version that would work on Linux and we did it in two months. Now hardly anyone uses the Windows Media version.

"We had been pushed into a corner by going for a download model rather than streaming," he says. But it turns out that the licensing restrictions on streamed content are much less problematic. This makes rights owners less demanding of digital rights management systems.

Butterworth says that platform neutrality is especially important as people increasingly expect to view content on a wide variety of devices. Bringing out his iPod Touch, he shows it streaming a

programme over a Wi-Fi internet connection. Devices such as the PlayStation Portable and the Asus EEE PC are ideal for watching TV on the move.

"People want to be able to watch content on anything and not be tied to one vendor's system," says Butterworth.



› **The *iPlayer* is far from ideal, but there's a will at the BBC to make it available to more users (and save us all money).**

› than H.264, which has a big barrier to entry in terms of unnecessary complexity. Unfortunately, the design of H.264 involved a lot of different academics, universities and companies, which meant that it was a case of design by committee. The commercial players were trying hard to insert their own proprietary intellectual property in order that they could participate in the licensing revenue that it generates. We didn't like that approach, so Dirac instead involved just a small team. We kept it as simple as we could."

That's not to say that the BBC doesn't have any patents on Dirac. It has a number of patents pending and says it's likely to apply for more, believing that this will protect the Beeb from claims from other codec designers. However, the BBC has granted a royalty-free licence for the use of those patents within the Dirac software.

The software that implements the codec is triple licensed. The main licence is the Mozilla Public Licence, which the BBC says "requires contributors not to enforce patent claims related to their contributions against each other". They

have also licensed it under the more standard GPL and Lesser GNU Public Licence (LGPL). The last of these is expected to be used by manufacturers.

Borer now hopes that other people will come up with new uses for the codec. "Dirac is not just about being free in terms of money; we're also giving people the ability to customise it to their own purposes. Because it's open source and patent-free, developers can do what they like with it without worrying about a licensing authority coming after them."

As an example, he explains that it became apparent that it would be useful to encode a separate stream that could hold additional data – such as signing for the deaf or subtitles – which could be kept apart from the main stream and switched on and off as

needed. Features in a similar vein that might not be immediately obvious to the original design team can be added by other users and fed back into Dirac's ecosystem. "When you come up with a very specific application that has a small user base, you haven't got a huge revenue stream to pay royalties and get permission," explains Borer. "Open source makes it easy for us to provide for lots of those sorts of uses." The Dirac team is already in talks with other European broadcasters that have been watching the project's development with considerable interest.


Box clever

"In principle, the Dirac codec could end up being used in Freeview boxes," says Borer. "While we haven't standardised it for Freeview yet, there are moves in the broadcasting world towards a next-generation box for digital broadcasting and, potentially, Dirac might be used in that." Freeview box manufacturers, Borer says, complain that too much of the cost of the boxes goes on royalties.

The BBC is about to start using Dirac in a very practical way. Much of the company's wiring that connects studios is designed for standard-definition television. With the introduction of high-definition, the BBC has faced the cost of replacing that infrastructure. So the Dirac team have partnered with Numedia Technology to produce a small pair of hardware units that compress and uncompress high-definition television so that it can fit down the BBC's current infrastructure and still be suitable for broadcasting. It is estimated that this hardware will save the corporation £11m, more than paying for Dirac's research and development and saving the BBC (and the licence payers) a considerable amount at the same time.

TIM BORER ON DIRAC


"It's not just about being free in terms of money; we're giving people the ability to customise it."



› **An Ingegria system at work on the set of *EastEnders*.**



› Using *Ingex* in the production control room while filming children's show *Bamzooki*.



› Tim Borer has worked on open source technologies such as subtitling software and video compression.



› The Linux-based *Ingex* system being tested in the Kingswood Warren lab.

Dirac is not the only BBC project that might find a use in next-generation Freeview boxes. Elsewhere in Kingswood Warren, Jeff Hunter and his team of engineers are researching how a high-definition, interactive Freeview box should work. Although the BBC is not in the business of selling set-top boxes, it has a lot at stake here. It is a major shareholder in DTV Services, the company that runs Freeview, and has licences to broadcast on two of Freeview's 'multiplexes,' and it's important to the BBC that future set-top boxes are able to provide the features that it wants to use.

Indeed, one of the problems with driving forward a platform like Freeview is that manufacturers aren't keen to implement new features unless all the other manufacturers are doing likewise – after all, it's a market where price is king, with Freeview boxes available in supermarkets for less than £30. If the BBC simply goes round asking companies to add things, it doesn't get a good response. Hunter's team have been working on a Linux-based set-top box that they plan to demonstrate to manufacturers. The prototype uses a standard x86 processor in a media centre case. Linux is an obvious choice because the box can freely build upon existing code. Of course, Linux is already being used in the set-top box market, particularly at the high end (in Sony Blu-ray players, for example). The BBC's work at the Freeview end of the market might well make the operating system irresistible.

The broadcasting of Freeview itself is also reliant on Linux. "Some of the back-end of the Freeview service runs on Linux," Hunter explains. "This includes key processes, some of which are mission-critical." Notably, Freeview's digital text service, the successor to Ceefax, is running on Linux. Outsourced to Red Bee Broadcasting Dataservices, it operates off 60 Linux machines running Red Hat Enterprise Linux 4, which are responsible for "building up the carousel of content and then providing real-time updates" as breaking news occurs, for example. Linux meets the BBC's needs because it's essential that the text service is reliable and does not go offline.

Second nature

"Linux use will increase at the BBC," according to Stuart Cunningham. "But people won't really notice it." He points out that although some of the Sony television cameras the BBC is buying are running embedded Linux, there's no big sign on them to let people know that they are Linux-based.

"We've always encouraged open standards," says Cunningham. "Encouraging manufacturers to make their equipment so that it's compatible with everything else we buy is better for the BBC. It lets us get better pricing because we can buy cheap commodity equipment. It's second nature at the BBC to favour open source and open technologies." **LXP**