



RawFlow

— using P2P to create virtual “multicasts”

Ingjerd Straand Jevnaker

Rawflow

RawFlow is one of several emerging technologies for streaming audio and video content over P2P networks. The Danish public-service broadcaster DR has been using RawFlow for streaming its online radio since 2003 and will shortly start streaming video with this technology. TSR, the Swiss national broadcaster, trialled RawFlow video streaming during the Tour de Suisse cycling race earlier this year and the Estonian public broadcaster, Eesti Radio, is also experimenting with it.

A broad overview of RawFlow is given in this article along with some results from its use by DR over the last three years.

As teenagers spend more time on the Internet than with traditional broadcast media, broadcasters need to move more of their content online. For example, in June alone, more than 2.5 billion videos were hosted on the social-networking site *YouTube*, equating to 200 terabytes of data per day (*Reuters: 16/07/2006*). And with 60% of all American teenagers having a personal profile on the social networking site *MySpace*, it is clear that broadcasters cannot ignore forever the effect of the Internet as a rich media platform.

However, increasing demand for rich media online has created a bandwidth problem for broadcasters. This so-called “bandwidth crunch” can be resolved, surprisingly enough, by the industry’s worst “foe” – ***Peer-to-Peer (P2P) technologies***. London-based technology company RawFlow has developed Intelligent Content Distribution (ICD) – a technology that creates a layer-7 (software-based) multicast which introduces scalability to internet broadcasting. This article aims to provide a technical overview of the P2P streaming landscape and RawFlow’s ICD technology platform.

The Internet as a “consumer friendly” media platform

The main digital alternatives to internet broadcasting – IPTV and DAB – have not quite picked up the speed that the broadcasting and teleco industries had expected. Yet, both remain niche standards even after years of deployment in Europe. The industry has failed to communicate the benefits of these platforms to consumers and, quite frankly in our view, what are they? IPTV offers essentially the same product and programming that digital cable and satellite providers have done for years, other than a different pricing model and limited Video-on-Demand (VoD) features. DAB deployment has been delayed or postponed in many countries and the DAB receiving devices are still too expensive. Other than in the UK and a few other European countries, the technology and broadcast content have not been compelling enough to make the consumer fork out a small fortune for a DAB receiver.

On the other hand, Internet broadcast is an open and democratic form of broadcast over *existing infrastructure* that can be accessed across platforms and devices. The Internet facilitates a distributed and collaborative environment for media production, where anyone can be a broadcaster. The future of broadcast media is *a la carte* content consumption where *Google* and other online aggregators are the new remote controls.

P2P - from foe to friend

The Internet is primarily built to handle point-to-point broadcast and problems arise when you want to broadcast point-to-many simultaneously – as the size of the audience participating in a live Internet broadcast increases, the bandwidth required for that broadcast increases proportionally. The traditional way to handle this cost is by limitation of the audience size or reduction of the quality, neither of which are ideal from a broadcaster's perspective.

In order for internet broadcasting to become scalable, a distributed approach must be employed. One such approach is **P2P streaming** – a technique much feared in the media industry – but nonetheless, in the quest for scalable internet broadcasting, P2P may become the media industry's best new friend. P2P technology is more than casually downloading the new Razorlight album for free over a piracy network. Despite being built on the same kind of technology, peer-based streaming should not be confused with p2p *file-sharing* and copyright infringement. A P2P computer network is a network that relies on the computing power and bandwidth of the participants in the network rather than concentrating it in a relatively low number of media servers.

Via P2P streaming, the bandwidth requirement of the broadcast is intelligently distributed over the entire network of participants, instead of being centralized at the broadcast's origin. As the audience grows, so do the network resources available to distribute that broadcast without adding any additional bandwidth costs. P2P technologies take the content the last mile to the consumers and help improve the consumer experience over relatively restricted DSL connections. Thus, a P2P architectural approach can be used to create a "virtual multicast".

Intelligent Content Distribution

RawFlow's *Intelligent Content Delivery* (ICD) technology is a live P2P-based streaming technology which enables broadcasters, content owners, content delivery networks and corporations to scale their webcasts to serve large audiences with reduced bandwidth and infrastructure costs. The ICD platform uses **grid technology** to distribute live video and audio streams. The grid network utilises the resources of its users to relay the stream to other users. Using this approach, as the audience grows so do the resources available to distribute the stream.

The RawFlow ICD Server is a piece of software that can be installed on a standard media server. This software provides the initial contact point for clients in the network. The ICD Server connects to the broadcaster's media and begins receiving the stream, maintaining a buffer of stream data in memory at all times. When a user clicks on a link to play the media, he or she will be asked to install a browser plug-in called the ICD Client. The ICD Client monitors the quality of the stream it is receiving and, upon any reduction of quality or loss of connection, it again searches the grid for available resources while continuing to serve the client's media player from its buffer.

The buffer prevents interruption to the stream and ensures that the end-user experience is not affected. The ICD Client first contacts the ICD Server and begins receiving the stream from it. If it

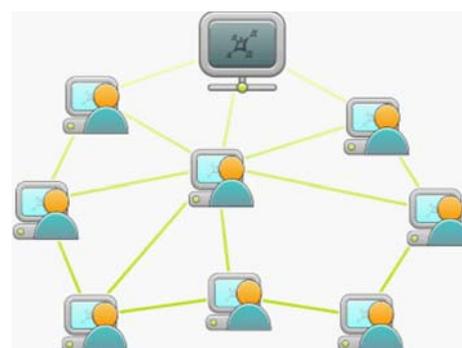
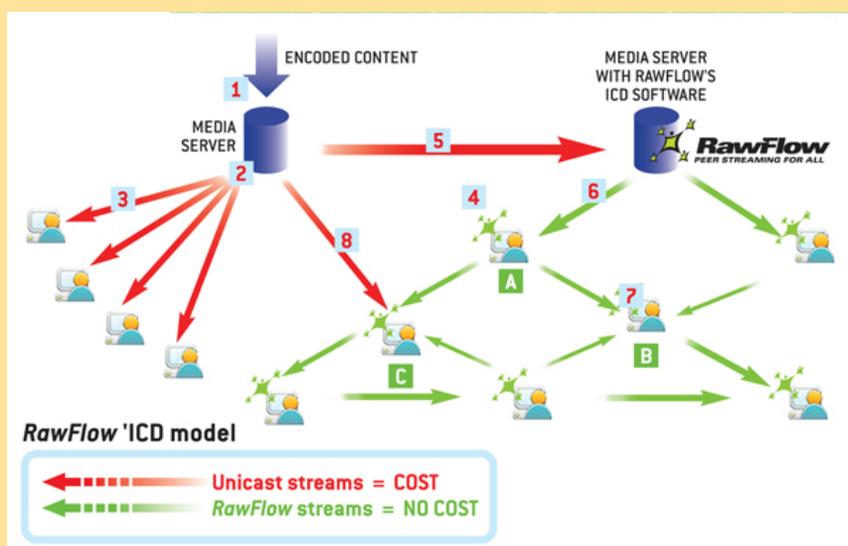


Figure 1
ICD network grid topology

How does Intelligent Content Distribution (ICD) work?

P2P technology utilises the redundant (and therefore free) upstream capacity of the end user. While the viewer watches video on the downstream, the video is relayed to a peer on the upstream to a grid of users.

- 1) Encoded content is sent to the media server.
- 2) The media server stores and streams content to the end user.
- 3) Each user receives a unique stream (unicasting). This results in cumulative bandwidth consumption equal to the number of users multiplied by the connection speed.
- 4) With RawFlow, when a user clicks to get the content, a connection is made to a media server running RawFlow's ICD software. The software manages the peering network.
- 5) The ICD "server" pulls a single stream from the media server.
- 6) The ICD "server" streams the content to the user (A).
- 7) The second user (B) joins the broadcast and receives the stream on A's redundant upstream. Multiple users can jointly contribute bandwidth so that every spare kilobit of bandwidth is used most efficiently.
- 8) When there is insufficient bandwidth available on the network to support the broadcast, the ICD software instructs the media server to deliver a "booster" stream to the network (C).



RawFlow Intelligent Content Distribution model

has available resources, the ICD Client also accepts connections from other clients in the grid to which it may relay a portion of or the whole stream it receives as requested.

Case study: Danish Radio via P2P

RawFlow works with a number of broadcasters worldwide, delivering video and audio securely via our Intelligent Content Delivery grid. RawFlow has been working with the Danish Broadcasting Corporation - DR - since 2003. DR is Denmark's oldest and largest media enterprise. DR runs four nationwide FM radio stations, thirteen Digital Audio Broadcast (DAB) stations and eleven additional internet radio stations, plus two major television channels and an extensive web portal, a version of which is accessible via mobile phone. DR is also Scandinavia's largest provider of online radio services.

RawFlow's ICD technology was adopted by DR to broadcast their most popular online radio channels in order to save the bandwidth costs associated with online services. In fact, although online TV may have greater challenges than online radio in terms of the quantity of data transmitted per

second, consumer behaviour means that more data in *total* may actually be transmitted by an online radio station, because radio can run in the background while you're working and not actively listening. All the DR audio streams are broadcast in low, medium and high quality (32, 64 and 128 kbit/s) using the Windows Media 10 codec.

RawFlow's technology is deployed so that the first listeners are connected directly to DR's own server cluster. The rest of the listeners are connected to each other and receive exactly the same broadcast as those connected directly to the source. They receive the audio stream from each other via their upload capacity. This means that DR does not have to use any further bandwidth to accommodate new users.

The peer solution is embedded seamlessly into the DR web player so that end users will not see or hear any difference. The only alteration to the user experience is that when a new listener accesses the stream for the the first time, he or she will be prompted after 10 seconds to accept the installation of a browser plug-in. This was thought to be an obstacle at first due to firewalls and internet security settings in corporate LANs but, in fact, more than 90% of the users of DR have accepted the plug-in.

DR: key statistics

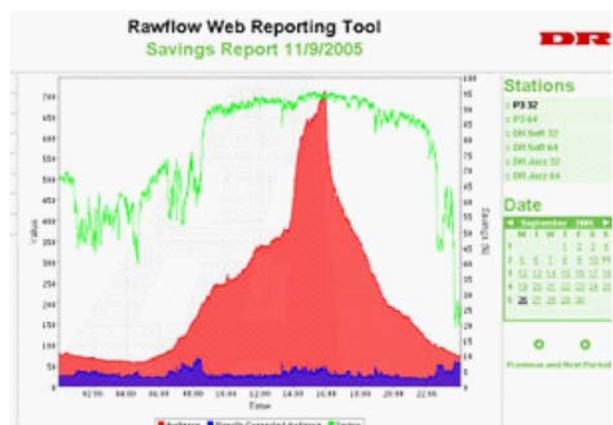
- 90% end-user acceptance
- 714,950 client downloads
- 100% server uptime
- Average Rawflow efficiency for DR was 72.9% of total traffic
- Peak Rawflow efficiency 99% of audience
- Five-fold increase in media server capacity
- More than 40 Million hours streamed through Rawflow ICD grid



Figure 3
A user is requested to install the browser plug-in, here white-labelled for DR
(courtesy of Danish Broadcasting Corporation)

Nearly as important perhaps, is that DR now has a complete reporting tool available, and is able to monitor exactly how many listeners they have at any point of time. They also get an overview of how many listeners are directly connected and how many are “peers”. The rate of efficiency (= savings) of P2P technologies is calculated based on the percentage of the audience who are “peering” as opposed to those receiving the media direct from the source (unicasting), which essentially means the percentage of “free users” relying solely on upstream/peering to get their content feed. At the time of the September 2005 screenshot (see Fig. 4) taken from the RawFlow statistics tool, the upstream capacity in Denmark was less than 256 kbit/s average upload. Yet, for DR, the *average* efficiency rate has been 72.9% and close to 99% for peak times. More importantly, that also means that DR has been saving between 72.9% and 99% of their bandwidth costs.

Figure 4
Daily Savings Report, Channel P3
32 kbit/s, 9 September 2005
The red area is the P2P audience and the blue area is the directly-connected audience. The green graph indicates the efficiency in %.
(courtesy of Danish Broadcasting Corporation)



More information about DR can be found at <http://www.dr.dk/>, one of the most visited portals in Denmark. You can also listen to the online radio channels from this location.

How efficient is P2P?

As peer-enabled streaming relies heavily on the upload capacity in the network, there is a limit on streaming efficiency, as long as the bandwidth is low. Additionally, the upstream capacity is very different from country to country.



Figure 5
Weekly Savings Report, Channel P3, 128 kbit/s, week beginning 11 September 2006
 The red area is the P2P audience and the blue area is the directly connected audience. The green graph indicates the efficiency in %. Typically for online radio the web-audience drops during the weekend.
(courtesy of Danish Broadcasting Corporation)

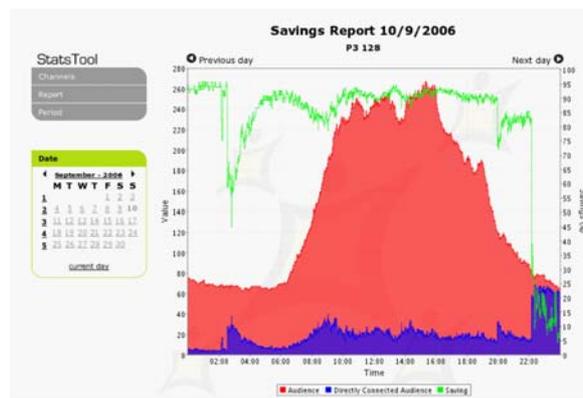


Figure 6
Daily Savings report, Channel P3, 128 kbit/s, 10 September 2006
 The red area is the P2P audience and the blue area is the directly connected audience. The green graph indicates the efficiency in %
(courtesy of Danish Broadcasting Corporation)

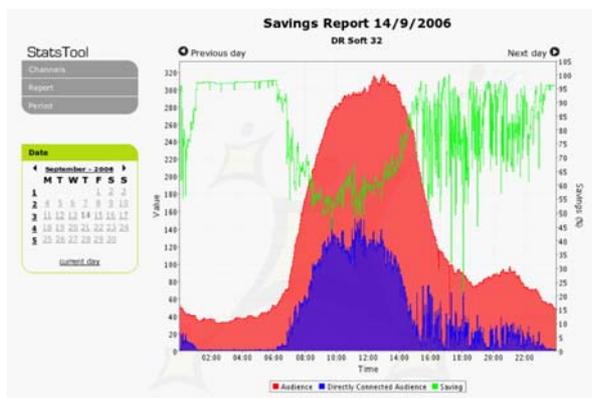


Figure 7
Daily Savings report, Channel DR SOFT, 32 kbit/s, 14 September 2006
 The red area is the P2P audience and the blue area is the directly connected audience. The green graph indicates the efficiency in %
(courtesy of Danish Broadcasting Corporation)

One of the challenges of P2P is that the only real way of testing the efficiency is in “live internet space”, i.e. it cannot be measured properly in a lab setting. Theoretically, efficiencies can be up to 99% but, as the internet is dynamic and changes by the second, it will vary according to the audience mix (users will have different ISPs / upload capacities). The efficiency will also be lower if the number of users drops below a certain level. RawFlow’s experience-based¹ efficiency is approximately 80% for a 128 kbit/s stream but the efficiency drops at higher bitrates due to upstream capacity limitations and other factors.

1. This number is based on historical statistics taken from ICD implementations in a number of countries, going back to 2002. The later implementations have shown higher efficiencies, also at higher bitrates. RawFlow is currently about to release version 5 of its ICD technology.

Efficiencies can be calculated based on the generic formula:

- Data rate < upstream capacity = more efficiency
- Data rate > upstream capacity = less efficiency

The RawFlow statistics tool provides an overview of the stream efficiency on a daily, weekly or monthly basis – one stream at a time. The efficiency will vary on a daily basis and is highly dependent on variations in the audience.

Figs 5, 6 and 7 show some screenshots taken from the RawFlow statistics tool.

As of yet, the only way to measure the bandwidth savings is by testing P2P in the “real-world” internet. However, RawFlow is looking at ways to develop new tools for measuring the efficiencies more accurately in a test-lab setting.

Network optimization

Network operators are under pressure to cope with growing traffic volumes. By combining Content Delivery Networks (CDNs) with RawFlow's live P2P technology, clients that use CDNs will benefit from lower broadcasting costs, greater scalability, improved broadcast quality, lower risk, increased capacity and global reach. Compared to traditional streaming, the RawFlow ICD improves network efficiency without increasing traffic within the network and without saturating the users' upstream bandwidth. Additionally, RawFlow has partnered with CacheLogic – a provider of traffic management, media delivery and network intelligence solutions – in order to reduce the traffic load on ISPs and to provide additional capacity when the upstream capacity is low.

CacheLogic has developed *VelociX* – a global network of P2P caches – which extends the low-cost, large-object delivery and high-scalability capabilities of P2P with the high-speed delivery and resilience of a globally-distributed network. RawFlow, with the assistance of CacheLogic, has extended its ICD technology to support CacheLogic's *Cache Discovery Protocol (CDP)* which enables the ICD Client to communicate with CacheLogic's *VelociX Media Delivery Platform*. The *VelociX* platform acts as regional distribution points to improve the network efficiency and reduces the load of P2P traffic on ISPs' networks. By limiting the number of feeds coming into an ISP's network, significant savings can be made.

The integration of CDP into the RawFlow P2P Client creates a dynamic peer environment where the *VelociX* platform can step in and provide additional capacity when needed, thereby ensuring that the user experience is maintained. This is particularly important when the average upstream capacity in a regional area is low.

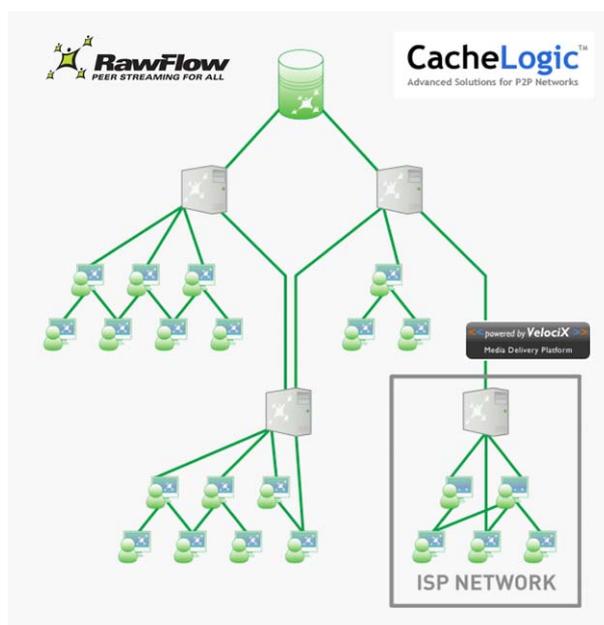


Figure 8
CacheLogic's *VelociX* Media Delivery Platform reduces the load on an ISP network and can provide additional capacity in regional areas where the average upstream capacity is low

RawFlow's business model

RawFlow's business model is to be a technology provider – but not to compete with webcasters or content delivery networks. Essentially, RawFlow provides the P2P delivery software and the user interface, i.e. the media player. The ICD software can be hosted on RawFlow's media servers but the company strategically does not provide content hosting or unicasting.

The ICD software is a flexible technology that can be implemented and integrated easily on top of any kind of infrastructure. Because the technology is platform-, format- and hardware-agnostic, it can fit into a variety of industries. Currently, RawFlow's technology is marketed to four main stakeholder groups:

- Broadcasters – the online media/entertainment sector (TV, radio and others);
- Corporates – universities, banks, consultancies and large global corporations;
- Networks – Content Delivery Networks and Internet Service Providers;
- “Prosumers” – online communities and early adopters of technology in the general public who are self-casting.

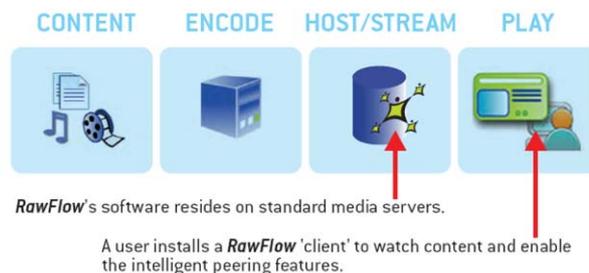


Figure 9
Where does RawFlow fit into the streaming value chain?

The pricing model and the offering vary according to the market. Based on a consultative sales process, the pricing models are customised and adapted to the user's needs. In general, broadcasters are offered an annual ICD software licence fee, where the charge is based on a combination of savings projections and data traffic through the network. RawFlow also offer a 30-day free trial of the ICD software.

Future outlook

While online TV at the moment may mainly consist of sharing low-quality video clips and user-generated content, the big question is: *what will TV on the Net mean in the future?* The use of online video content is growing rapidly and this is creating an increasing amount of last-mile infrastructure challenges. Therefore, in the near future, the legal and commercial P2P landscape will be further developed, and we can expect to see more companies coming up with new peer-based delivery solutions. We can also expect that content delivery networks will peer-enable their networks in order to be able to offer broadcasters the ability to reach more consumers. P2P networks are dependent on traditional unicasting as the fallback and initiator of the streaming. RawFlow will aim to partner with CDNs in the near future.

Combined with increased bandwidth, it seems likely that peer-enabled streaming technologies are more than a quick-fix to the bandwidth crunch while we're waiting for "real" multicasting (the way in which the market is going forward, it looks like it will be quite a long wait). Online broadcasting can only be truly viable by employing a distributed approach, and P2P media delivery *can* support broadcasters, content owners and music labels today. So, in many ways, you could say that P2P has gone from "foe to friend" for the media and entertainment sector.



Ingjerd Jevnaker works as Marketing Manager for P2P streaming company Raw-Flow Inc. Previously she worked as Account manager for Ashridge New Media, a London-based web design agency, and as Marketing Manager for Spreadshirt GmbH, a German online merchandising company, where she was in charge of online marketing activities and business development for the Scandinavian market. Additionally, she has worked as Research Assistant at the Norwegian School of Management, department of Innovation and Economic Organization, and for a Boston-based biotechnology start-up; 3DM Inc.

Ingjerd Jevnaker holds an MSc. in Business Administration and a BSc. in Marketing and Management. On a personal note, she is an active blogger and has also made presentations to the media and music industries in the UK and elsewhere on the topics of digital futures and the challenges of convergence.

Further information

RawFlow

E-mail: sales@rawflow.com

Tel: +44 (0)20 7480 4220

Web: <http://www.rawflow.com>

CacheLogic

e-mail: sales@cachelogic.com

Tel: +44 (0)12 2339 5914

Web: <http://www.cachelogic.com>
