



ICRA's experience of the technical and policy issues related to content labelling.

1 Abstract

ICRA's descriptive vocabulary allows content providers to label their sites in as objective a way as possible. At present this is done using PICS labels that are easy to use in simple websites but that pose significant problems for larger web properties. The limitations of the PICS system presents both policy and technical problems for the largest suites of websites. A new solution based on RDF seems logical, however this approach itself is not without drawbacks. There is as yet no method for linking multiple resources with a single description let alone a method that is easy to use, flexible and attractive. A possible solution proposed here uses a similar concept to the Policy Reference File defined under P3P and links content rating with many other forms of metadata to attract the widest possible support. Such a method would have equal applicability to fixed and mobile internet services as well as a full range of other digital media.

2 Introduction - a brief history of ICRA

The Internet Content Rating Association was created in 1999 under a project that had the backing and direct involvement of many of the biggest names on the internet, as well as the European Union's Safer Internet Action Plan. The aim was relatively simple if ambitious: to develop a system that would allow webmasters to describe their content using an internationally acceptable, cross-cultural language. Parents and other carers would then be empowered to choose which types of content they did and did not want their children to have access to, reflecting their own values rather than those of any one else.

This was not a new idea. Self-labelling had been around since the mid 1990s with the advent of PICS¹ and the early rating vocabularies, notably RSACi which ICRA superseded. The important difference about ICRA is that the descriptive vocabulary is binary in nature – an element is present or absent – so that subjective ideas like a “scale of nudity” were avoided. The vocabulary was developed by an international panel of experts and published in December 2000. The ICRA website now carries the descriptor set in 6 primary

ICRA's experience of the technical and policy issues related to content labelling.

languages, including Chinese and the major European languages, with translations available in a further 4 languages.

Approximately 100 webmasters a day visit the ICRA website, fill in the questionnaire, and generate a PICS label that they can then add to their content. Filters can then read those labels and take action based on the description found and the parental settings.

Parents need not necessarily work with the ICRA vocabulary directly. Third party organisations may offer their interpretation of it, perhaps an age-based classification or one that reflects a particular life-view. These could be encoded in what we call "templates" (but are actually PICSRules Files²) that can be imported into filters. This greatly simplifies the parent's job while the providers benefit from a single international system that does not impose the value-judgements of any one country or cultural philosophy.

An initial aim for ICRA was that Microsoft would update the Content Advisor function in Internet Explorer to read ICRA labels rather than the old RSACi system. Indeed, Microsoft was to be a partner in the original project. That proved to be impossible, however, as EU rules prevented the US-based software company from being an official project partner. By early 2001 it had become clear that ICRA would need to offer an alternative label-reading system.

The result was ICRA *filter*, a tool that demonstrated the concept of filtering against ICRA labels. By the time of its release in March 2002, work had already begun on an enhanced filtering system now known as ICRA *plus*³. Again co-funded by the EU's Safer Internet Action Plan, the project brought ICRA together with the Software Knowledge Engineering lab at NCSR "Demokritos" and Spanish filtering company Optenet.

Like ICRA *filter*, ICRA *plus* is a free tool that can block or allow access to sites carrying ICRA labels by matching parental choices against the types of content declared in the label. The big difference, however, is that ICRA *plus* offers users enhanced functionality by working with optional extra filters. Whatever technology the optional filters use, their result is fed into ICRA *plus* in the form of an ICRA label, thus using the internationally-respected ICRA vocabulary to deliver the description of a website to the parent, even if the site in question is not labelled itself.

ICRA stands by its commitment to self-labelling as the best, most open and democratic method by which a provider can describe his or her content. Alongside self-labelling, however, the same vocabulary can be used by other technologies in the international, cross-cultural, multi-lingual medium of the internet.

It is this multi-faceted approach that is driving ICRA forward and allowing the organisation to find relevance in new areas.

3 Ideologically pure, sometimes practically difficult

In theory, self-labelling using PICS is easy. Servers are configured to include labels in the HTTP response headers or http-equivalent meta tags are added to HTML pages. For websites that have a single webmaster and run from a single server, this is easily achieved. For larger web “properties” the proposition is very different. There are two principal barriers to labelling the largest of sites:

1. There is no such thing as “the webmaster”, or even a senior webmaster who has overall responsibility for the site. There are divisions, there are product managers, marketers, policy advisors, production staff and, occasionally, engineers. Who is responsible for labelling in such an organisation?
2. “Web pages” are created by content management systems that may derive content from any number of servers in any number of locations. Where would the labels be added? Which domains are relevant to a particular ‘page’? How does one define ‘page’ given that HTTP is a stateless protocol?

If all your content should be labelled the same way, these problems become manageable. You can put in place policy and technical systems to ensure that all your content carries the same label. In practice this is rare although British Telecom is perhaps a good example. Their label is:

```
pics-1.1 "http://www.icra.org/ratingsv02.html" 1
gen true for "http://www.btopenworld.com" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://static.btopenworld.com" r (nz 1 vz 1 lz 1 oz 1
cz 1)
gen true for "http://213.121.143.196" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://www.btinternet.com" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://btopenworld.com" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://btinternet.com" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://www.btyahoo.com" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://btyahoo.com" r (nz 1 vz 1 lz 1 oz 1 cz 1)
gen true for "http://btyahoo.co.uk" r (nz 1 vz 1 lz 1 oz 1 cz 1)
```

This lengthy label list covers the 9 domain names BT uses for its primary websites. This label had to be inserted into no less than 75 page templates. Manually⁴.

The problem becomes even more complicated if there are sections of the site that should be labelled differently. PICS supports this, yes, but the task of adding suitable PICS labels in the correct places becomes a logistical headache. Furthermore, it is common amongst the largest online properties to include user-generated content or content supplied under contract by third parties. These content types present problems for the policy team – how do you label content you don't control?

ICRA's experience of the technical and policy issues related to content labelling.

3.1 Other problems with PICS

As well as the policy problems presented by user-generated or third party content, and the sheer technical headache of including a label in all the places it needs to go, there are other problems with PICS labels too. The label on the ICRA website is as follows⁵:

```
pics-1.1 "http://www.icra.org/ratingsv02.html" l
gen true for "$ENV{HTTP_HOST}" r (cz 1 lz 1 nz 1 oz 1 vz 1)
gen true for "http://a.tribalfusion.com" r (cz 1 lz 1 nz 1 oz 1 vz 1)
gen true for "http://m.tribalfusion.com" r (cz 1 lz 1 nz 1 oz 1 vz 1)
gen true for "http://ad.doubleclick.net" r (cz 1 lz 1 nz 1 oz 1 vz 1)
gen true for "http://m2.doubleclick.com" r (cz 1 lz 1 nz 1 oz 1 vz 1)
gen true for "http://m.doubleclick.net" r (cz 1 lz 1 nz 1 oz 1 vz 1)
gen true for "http://tribalfusion.speedera.net" r (cz 1 lz 1 nz 1 oz
1 vz 1)
gen true for "http://pagead.googlesyndication.com" r (cz 1 lz 1 nz 1
oz 1 vz 1)
gen true for "http://pagead2.googlesyndication.com" r (cz 1 lz 1 nz 1
oz 1 vz 1)
gen true for "http://view.atdmt.com" r (cz 1 lz 1 nz 1 oz 1 vz 1)
```

The first label writes its domain in dynamically, so that whichever of ICRA's 8 domain names is used, the label is valid. There then follows a series of labels for the domains used by ICRA's advertising agency. Ideally, of course, the banner advertisements would arrive labelled themselves but this isn't the case and, politically, ICRA needs to offer a website that is fully labelled. The only way this can be done is to include labels for the domains used by the agency.

3.2 More is possible

It is possible to label large sites using PICS labels. There are two prime examples of this in Germany, broadcaster RTL and major ISP, T-Online.

RTL Deutschland labels its online content through configuring its server to include labels in HTTP response headers thus⁶:

```
HTTP/1.1 200 OK
Date: Wed, 21 Apr 2004 09:54:22 GMT
Server: Apache
pics-label: (pics-1.1 "http://www.icra.org/ratingsv02.html" l r (cz 1
lz 1 nz 1 oz 1 vz 1) gen true for "ad.de.doubleclick.net/adj" r (cz 1
lz 1 nz 1 oz 1 vz 1) gen true for "ivwbox.de" r (cz 1 lz 1 nz 1 oz 1
vz 1) gen true for "count.rtl.de" r (cz 1 lz 1 nz 1 oz 1 vz 1)
"http://www.rsac.org/ratingsv01.html" l r (n 0 s 0 v 0 l 0) gen true
for "ad.de.doubleclick.net/adj" r (n 0 s 0 v 0 l 0) gen true for
"ivwbox.de" r (n 0 s 0 v 0 l 0) gen true for "count.rtl.de" r (n 0 s
0 v 0 l 0))
Last-Modified: Wed, 21 Apr 2004 09:47:18 GMT
ETag: "1a700c-4e9d-40864326"
Accept-Ranges: bytes
Content-Length: 20125
Connection: close
Content-Type: text/html
```

Different labels are served in the erotic section of the RTL.de website.

T-Online includes a lengthy meta tag in its page templates thus⁷:

```
<meta http-equiv="pics-label" content='(pics-1.1
"http://www.icra.org/ratingsv02.html" 1 gen true for "http://www.t-
online.de" r (nz 1 vz 1 lz 1 oz 1 cz 1) gen true for
"http://toi.ivwbox.de" r (nz 1 vz 1 lz 1 oz 1 cz 1) gen true for
"http://count.passul.t-online.de" r (nz 1 vz 1 lz 1 oz 1 cz 1) gen
true for "http://banner.t-online.de" r (nz 1 vz 1 lz 1 oz 1 cz 1) gen
true for "http://pix.t-online.de" r (nz 1 vz 1 lz 1 oz 1 cz 1) gen
true for "http://www3.service.t-online.de" r (nz 1 vz 1 lz 1 oz 1 cz
1) "http://www.rsac.org/ratingsv01.html" 1 gen true for
"http://www.t-online.de" r (n 0 s 0 v 0 l 0) gen true for
"http://toi.ivwbox.de" r (n 0 s 0 v 0 l 0) gen true for
"http://count.passul.t-online.de" r (n 0 s 0 v 0 l 0) gen true for
"http://banner.t-online.de" r (n 0 s 0 v 0 l 0) gen true for
"http://pix.t-online.de" r (n 0 s 0 v 0 l 0) gen true for
"http://www3.service.t-online.de" r (n 0 s 0 v 0 l 0))' />
```

Why should RTL and T-Online go to so much trouble to do this? T-Online has long been a supporter of ICRA, indeed was a founder member, so there is high level policy support but there's more to it than that. In April 2002, a new federal law, the Jugendmedienschutz-Staatsvertrag, came into force. This puts the onus on content providers to protect children from potentially harmful material without prescribing in detail how this should be done. Self-labelling is seen as one possible route to this and therefore the benefit of labelling is more clear.

Other organisations have looked at labelling their web material to the same degree and have not been able to justify the considerable time and resources required. Several household-name websites are what ICRA terms "partially labelled". That is, there is a label for the primary domain, but the sub-domains, including those from which images originate for the homepages, are unlabelled⁸.

4 The cost-benefit analysis

The preceding discussion suggests that whilst there is a clear willingness to label and high-level policy support for it, the amount of effort that can be assigned to the task is limited. As with any other activity, managers must carry out some sort of cost-benefit analysis of labelling if resources are to be devoted to it. The table below attempts to summarise some of the issues.

ICRA's experience of the technical and policy issues related to content labelling.

Some of the costs

- Identification of personnel responsible for online content. This sounds easy but can be difficult in large organisations.
- Categorisation of online material. How should content be described?
- Technical understanding of PICS. A technical person needs to understand how PICS works and how it might fit into their existing architecture.
- Establishment of system to deploy and maintain the labels consistently
- Assignment/training of relevant staff

Some of the benefits

- Potential protection from censorious legislation
- Potential protection against litigation
- Protection of children
- Corporate responsibility
- PR opportunities
- Shield against some forms of negative comment

The cost-benefit analysis will produce a different result for different organisations in different circumstances, however, ICRA is seeking to tip the balance significantly by both increasing the benefits and reducing the costs.

5 Isolating the problem to find a solution

PICS is rarely credited with the flexibility and power it actually has. Even so, we have seen that it does have severe limitations that are compounded by the stateless nature of HTTP. One simple change would be to make the quoted string in a 'for' statement a true URI so that a label that said:

```
gen true for "http://example.org"
```

would also cover <http://www.example.org> and <http://subdomain.example.org>. But there are more promising solutions than trying to re-work PICS which was designed for one thing – content rating – and little else.

RDF⁹ offers the ability to encode not just content rating but the whole range of metadata. It is finding support in areas like government websites where vocabularies such as Dublin Core have most relevance. Even so, it is not yet a part of standard web design in the way that, for example, CSS has become. I suggest that two pieces are missing:

1. A method of adding RDF metadata that is flexible and as easy to use in a large, complex suite of web properties as it is on a personal website.
2. A positive incentive to add RDF metadata to digital content

Each of these is addressed below.

5.1 A flexible method of applying RDF

During the most recent round of revisions, the `aboutEach` prefix was dropped from RDF. This was done for good reason¹⁰ but as a result, an RDF description can only be applied to a single URI. In other words, not just each page on a website would need its own RDF description, each element, such as images, script files etc. would all need their own RDF description if the site were to be fully 'labelled'.

Discussions with large content providers make it clear that a system is required that will allow the machine equivalent of statements like:

1. `http://example.org/ratings/#rating1` applies to all content on this domain *except* resources A, B, and any URL on the C subdomain.
2. `http://example.org/ratings/#rating2` should be applied to the C subdomain

As RDF can carry the full range of metadata, one can imagine a more comprehensive set of statements:

1. All material on this domain is copyright Example Publishing
2. The privacy policy for all pages on this domain can be found at `http://privacy.example.org`
3. Description for domain: An example description
4. Publisher: Example publishing
5. Keywords for `http://example.org/animals`: birds, mammals, arthropods
6. `http://example.org/ratings/#rating1` applies to all content on this domain *except* resources A, B, and any URL on the C subdomain.
7. `http://example.org/ratings/#rating2` should be applied to the C subdomain

This has some similarities with the P3P reference file concept¹¹. Such a reference file could be pointed to in the usual ways, either by a LINK REL or an HTTP response header. The reference file might contain RDF data itself or point to descriptions elsewhere, the most important factor though is that the reference file and the descriptions themselves can be held separately from the content if so desired, and that a single description can be linked to, or specifically NOT linked to multiple resources.

5.2 Creating a positive incentive

The reasons for which providers may add ratings to their content have already been discussed but why does any webmaster add any non-essential element to their site?

- Privacy policies help users feel confident about giving personal information.
- Trust marks help users feel more confident about, for example, spending money with online retailers
- Keywords and descriptions can help others to classify the material

Different content providers will place differing importance on these elements but visibility in search engines is a near-universal desire. If the metadata, including trustmarks, privacy policies and content ratings available on a website were machine readable and, importantly, machine verifiable, they would be of greater potential use to search engines.

To extend the metadata reference file idea discussed above a little further, one can imagine an entry like:

- This domain has been awarded a quality assurance label by the Quality Label Agency <<http://qualitylabelagency.org>>. This can be verified at <http://qualitylabelagency.org?id=example.org>

This would allow a user agent to interrogate the Quality Label Agency's database and get a real-time confirmation that the certificate was valid. In such a scenario, a search engine might choose to add an icon next to a website's entry in the search results to show that a quality label was present. Real-time content analysis might allow similar visual cues to indicate the presence of descriptive metadata that was *probably* correct.

6 Tipping the cost-benefit analysis: Quatro

If the same basic technique can be used to add everything from privacy policies to descriptive metadata, perhaps even to link to style sheets and script files, its usefulness becomes obvious. From ICRA's point of view, the logistical problems of "getting content ratings into the system" are greatly reduced at a stroke. Furthermore, if that same technique allows search engines to locate and verify such metadata, it becomes useful to them.

There is little of greater interest to content providers than something that is potentially useful to search engines.

Led by UK publishing consultancy Pira International and ICRA, The Quality Assurance and Content Description Project (Quatro) is about to start work in this area. Quality labelling organisations, policy groups, academics and technologists from across Europe are coming together to look at the various issues. From a policy point of view, would consumers, eCommerce operators, parents, software manufacturers, mobile operators and content providers

benefit from such a multi-faceted system? From a technical point of view is a metadata reference file a workable solution? If not, what is? If so, is there a better one anyway?

As an XML-based technology, RDF can be deployed just as easily in mobile communications infrastructure as on the fixed internet, as well as any other medium that has occasional or permanent network access such as games consoles and digital TV. The potential is significant.

The Quatro project has an open communication channel with the work being undertaken by IA Japan and Keio University on content rating for mobile services. Members of the W3C's Semantic Web team are involved in both projects and ICRA is working to set up a North American arm for Quatro, such has been the level of interest from many different organisations.

7 Conclusion

Rating web content using PICS is possible and it works. Many content providers have labelled their sites and the IA Japan has created a large-scale system based on a PICS label bureau. However, self-labelling using PICS for the more complex content management system-driven sites is difficult both from a policy and technical point of view. A new system, almost certainly based on the maturing RDF Recommendations, has the potential to make a real impact. For different reasons, many organisations are examining what needs to be done and what can be done. Encouragingly, there seems to be consensus on at least the outline of the solution.

*Phil Archer
CTO, ICRA
April 2004*

8 References

1. <http://www.w3.org/TR/REC-PICS-labels>
2. <http://www.w3.org/TR/REC-PICSRules>
3. See <http://www.icra.org/icraplus/>
4. See, for example <http://www.btyahoo.co.uk/default/>.
5. See <http://www.icra.org>
6. <http://www.delorie.com/web/headers.cgi?url=http%3A%2F%2Fwww.rtl.de>
7. See <http://www.t-online.de>
8. The label tester at www.icra.org/label/tester/ offers a method of seeing ICRA labels on sites and how they are interpreted.
9. <http://www.w3.org/RDF/>
10. <http://www.w3.org/2000/03/rdf-tracking/#rdfms-abouteach>
11. http://www.w3.org/TR/P3P/#ref_file