

Building Trust for E-commerce: Collaborating Label Bureaus

Michael Shepherd, Anil Dhonde, Carolyn Watters
Web Information Filtering Lab
Faculty of Computer Science
Dalhousie University
Halifax, Nova Scotia, Canada B3H 1W5
{shepherd | dhonde | watters} @cs.dal.ca

Abstract

A major problem facing the full deployment of business-to-consumer (B2C) e-commerce is the development of trust on the side of the consumer. People develop trust in a business through their own previous experience with that business or through reports about that business from trusted third parties or other consumers. While a great deal of effort has been spent on privacy and secure transactions, resulting in seals of approvals and trustmarks, little has been done in the way of developing systems that will assure the consumer that the e-commerce site with which they are dealing is “legitimate” in such things as delivery, return policies, etc. A problem with seals and trustmarks is that they are an “all-or-nothing” approach and give a single value (the seal). A more flexible system is required that rates e-commerce sites along multiple dimensions and allows the consumer to determine which dimension(s) is of importance at that moment in time. These rating should be done by trusted third parties and stored in Label Bureaus. A prototype architecture has been developed that supports distributed Label Bureaus and multiple rating systems that rate along multiple dimensions. The architecture is based on a Label Engine and an Application Server that retrieves labels from one or more Label Bureaus and applies a rating algorithm to merge the various rating systems. An architecture that supports third party labeling allows the consumer to develop a rational trust level in unknown businesses.

1. Introduction

E-commerce on the Web is made up of two major categories of transactions; business-to-business (B2B) and business-to-consumer (B2C) [7]. While B2C transactions on the Web are growing, they are not growing as quickly as B2B transaction. Forrester Research expects that B2B e-commerce will be responsible for approximately \$1.5 trillion in transactions by 2003, 14 times larger than their estimate for B2C transactions [10]. By this estimate, B2C would represent only about 7% of such transactions by the year 2003. Other, more conservative estimates, have the B2C share of e-commerce falling from its current 27% of dollar transactions to 17% by the year 2002 [15]

There are many reasons for the slower growth in B2C than in B2B e-commerce, but one reason has to do with the much slower development in e-commerce supporting services (such as security and payment services) on which B2C e-commerce relies. These

supporting services are important in creating the “legitimacy” conditions required for trust to develop on the part of the consumer. Some important “legitimacy” conditions [15] that allow trust to develop include:

- the sellers are who they claim to be
- the seller has right of sale over the item in question
- the transaction and payment mechanisms are available, legal and secure
- information about the buyer is not redistributed to other organizations or used for other purposes than for which it was intended
- the item sold corresponds to its description and is suitable for its intended purpose
- the purchased item can and will be delivered to the buyer
- the buyers are who they claim to be
- the buyer has the resources to purchase the item

All but the last two items are issues of “trust” from the perspective of the consumer, while the last two items are issues of trust from the perspective of the seller. In this research, we have concentrated on the issue of trust from the perspective of the consumer. This is a difficult issue to deal with as it cannot be resolved by strictly algorithmic mechanisms. Current approaches include seals of approval and trust marks. The Better Business Bureau Online [3] has two programs, a privacy seal program and a reliability seal program. After evaluation by the BBBonline, a company can receive either or both seals that can be displayed on their Web site. TRUSTe [16] awards a seal or trustmark to an e-commerce site for display on their Web site if the company adheres to established privacy principles such as indicating to the user what information is being gathered and how it will be used, etc. Such seals are valid, of course, only at the time of the evaluation itself.

A major problem with seals is that a seal represents a summarization of a number of different dimensions along which that business has been evaluated and the consumer may not be aware of the dimensions, the scale of allowable values on each dimension, or the values that were assigned to individual dimensions. The consumer only sees that the business in question either has the seal or that it does not have the seal. At best, the seal may have different levels, but this is still a rating on a single dimension of approval. To interpret the seal, the consumer must check with the issuing authority to determine the date of issue and what the level means. An alternative approach used by eBay [4] is to gather feedback from both buyers and sellers on each other. Comments are gathered and buyers and sellers are awarded one positive point for each positive comment and a negative point for each negative comment. Both the points accumulated and the actual comments are available to everyone. Although it is an interesting and useful evaluation, it is not controlled and the evaluators are not necessarily trusted third-parties.

A better approach would be the use of labels where a label consists of a set of dimensions and the business has been rated and assigned a value on each of these dimensions. For example, a Web site might be rated by a rating service on the dimensions of violence, sex and language. To engender trust, such ratings must be done by trusted third parties and stored in independent label bureaus, as described below. Then, the consumer can decide

if they wish to do business with the target business not based on an overall evaluation of that business but by an evaluation of the scores along those dimensions that are of importance to that consumer for that transaction.

For example, a consumer may need a certain item immediately. Business A is rated poor for its time-to-delivery of purchased items but very high on security and privacy of both communications and the Web site itself. Business B is rated excellent on its time-to-delivery but poor on security. The choice is now the consumer's; how badly does the consumer need that item and how important is the issue of security to that consumer in this particular instance?

A problem arises, of course, if the label bureau has a label for one but not the other of these businesses. Also, the consumer may want a second opinion, i.e., the consumer may not be comfortable with a label from a single label bureau. In this case, multiple label bureaus would have to be consulted and their labels integrated in some fashion. A similar system is being developed for the rating of medical information [6]. In that scenario, doctors, medical societies and associations, will critically appraise Internet information and act as third-party raters of the information with respect to the value and trustworthiness of the information. Consumers and professionals may subscribe to a multitude of these services to get, automatically, ratings from different perspectives when retrieving information from the Internet.

The aim of this research is not the development of rating systems for e-commerce, rather it is the development of a generic architecture for collaborative label bureaus as described above. The remainder of this paper describes label bureaus and rating systems in general and how they can be used with PICS [9] and PICSRules [5] for building trust in B2C transactions. This notion is then expanded to describe the generic architecture that allows collaborative rating based on the combination of labels from independent label bureaus.

2. Label Bureaus

Label Bureaus were developed initially as trusted third-parties for the storage and distribution of labels to be used by Web browsers and applications to protect children, as much as possible, from *unintentional access* to possibly legal, but nonetheless objectionable material [14].

Labeling refers to schemes to assign content related labels to URL' s and/or specific Web pages. The individual rating protocols exist, in general, separate from products or applications using these ratings. These labels can be stored as part of the Web page or separately from the Web page in a database. Labels may be the result of self-rating (first-party), community rating by interested users (second party), or third-party authority rating [1]. Self-rating refers to the practice of Web page creators to provide a label describing the content of those pages. The quality of these labels reflects the self-interest of the producer. Community rating schemes are informal efforts by like-minded consumers of Web pages to share their knowledge of the content of Web pages with each other. Third-party authority rating refers to groups or businesses that take on the service

of providing labels that describe the content of Web pages for the use of customers or their community of users.

In this paper, we restrict our view to trusted third-party raters and the independent Label Bureaus in which these labels are stored. Figure 1 is an example of label bureaus and their use. The third-party raters examine a Web site and generate a label and store it in the Label Bureau. When users wish to access the Web site, the labels for that site are retrieved from the Label Bureau and may be used in making decisions regarding that site.

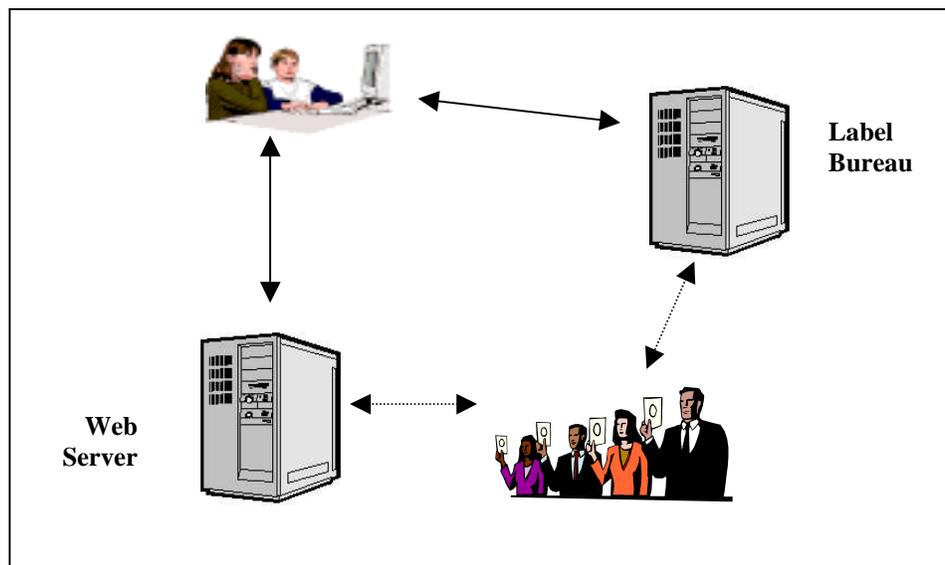


Figure 1. Label Bureau Access

Although different labeling protocols or schema are available, only a few are in wide use on the Internet and currently deal with the content of the Web site rather than the rating of the business behind the Web site. Nonetheless, the structure of these schema can be used in the e-commerce domain. Two examples of schema for Web filtering for children are:

- The RSACi system [12] rates sites on four categories (language, nudity, sex, and violence) and provides users with information on the levels (0 to 4) of these categories in given sites. RSACi is now used by Internet Explorer, CyberPatrol, and CompuServe (in the USA and in Europe).
- SafeSurf [13] is an eleven category rating scheme (age range, profanity, heterosexual themes, homosexual themes, nudity, violence, (sex, violence, and profanity), intolerance, glorifying drug use, other adult themes, gambling). Each

major category is subdivided into multiple subcategories. SafeSurf is used by both Netscape Navigator and Internet Explorer.

A layer cake model [1] for the combination of first and third-party ratings schemes has been proposed as part of the Self Regulation of Internet Content project [2], and has relevance to B2C e-commerce. As shown in Figure 2, such a model has three layers, all resting on a ‘plate’ of technologies such as PICS and XML for implementation of the model. The first layer of the cake is a basic vocabulary that will be used by first parties in rating their sites. The second layer of the cake consists of ratings templates created by third parties. These templates may reflect different legal systems, cultures, or interest groups. The third layer of the cake consists of a set of ratings of individual sites created by trusted third parties using the vocabularies and templates of levels one and two. This level also includes filtering software to be applied to these ratings. By separating the vocabulary elements from the construction of templates, we can better allow third parties to reflect their value systems while still preserving the vocabulary of the first-party raters.

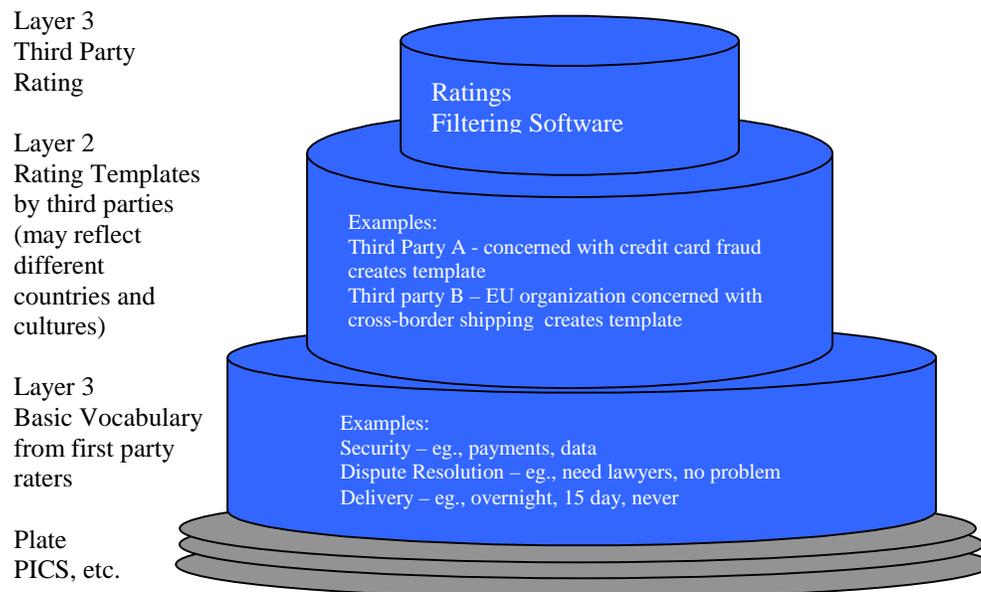


Figure 2. Layer Cake Model

3. PICS

Normally, PICS labels (Platform for Internet Content Selection [9]) are used to rate Web sites to protect children from unintentional access to objectionable material. In this research, PICS labels are used to rate e-commerce sites along multiple dimensions in an effort to create the legitimacy conditions for the development of trust on the part of the user.

The use of labels for site description requires several steps: creation of labels, storage of labels, maintenance of labels, dissemination of labels, authentication of label values, and finally the use of available labels by the filtering software to classify given entities as acceptable or not acceptable for the given context. The primary platform for associating labels with Web pages and specifying the format of these labels is PICS. PICSRules [5] form the primary platform for writing profiles for the filtering of Web pages based on these labels.

PICS specifies three methods by which these labels may be transmitted: in an HTML document using the META tag, with a document transported via a protocol that uses RFC-822 headers, and separately from the document where a client requests labels from a Label Bureau that runs the HTTP protocol. In our architecture, only this third method, i.e., requesting the label from a Label Bureau, is of interest.

The following is an example of a PICS label list, adapted from the W3C example, as it might be applied in an e-commerce situation. In this example, "<http://www.ecomm-rating-service.org/v1.5>" is the URL of the rating service that produced the label for the company represented by the site at <http://www.xyz-ecomm.com/>. The label was created by John Doe and is valid from November 5, 2000 until November 5, 2001. The label is providing ratings along three dimensions. Assuming that the ratings are on the scale of 0 through 10 where 0 is poor and 10 is excellent, then the rating for this e-commerce site was poor for security of payments and customer data, but was very good for fast delivery.

```
(PICS1-1 'http://www.ecomm-rating-service.org/v1.5'  
  by 'John Doe'  
  labels on '2000.11.05'  
  until '2001.11.05'  
  for http://www.xyz-ecomm.com/  
  ratings (secure_payment 2 secure_data 2 fast_delivery 8)  
)
```

The following example PICSRule could be applied to this label. The rule would alert the consumer before the transaction takes place if the company associated with this label is either rated very poorly with respect to secure payments or not rated highly enough for its ability to delivery in a timely fashion. The attribute "secure_data" is not pertinent to this rule.

```

(PicsRule-1.1
(
  serviceinfo (
    http://www.ecomm-rating-service.org/v1.5
    shortname 'ECR'
    bureauURL http://www.ecomm-label-bureau.org/ratings
    UseEmbedded 'N'
  )
  Policy (RejectIf '((ECR.secure_payment < 2) or (ECR.fast_delivery < 7))')
  Policy (AcceptIf 'otherwise')
)
)

```

In this example, labels embedded in documents are ignored (UseEmbedded 'N') and only labels retrieved from the Label Bureau site (<http://www.ecomm-label-bureau.org/ratings>) using the rating scheme '<http://www.ecomm-rating-service.org/v1.5>' are used to assess the companies.

4. Architecture

It was felt that, to build trust, it would be best to have multiple third-party ratings of individual companies. The third-party resolves the issue of self-interest found in first-party rating. Integrating multiple third-party ratings would give some protection against both overly positive and overly negative evaluations. In the architecture that has been developed, these multiple ratings are stored in multiple Label Bureaus and a single Label Engine maintains a database of Label Bureaus. Multiple labels are integrated by the Application Server. Thus, a client request for labels goes to the Label Engine, not to the individual Label Bureaus. This architecture provides the user with the following capabilities:

- **Rating.** Labels can be retrieved from various Label Bureaus regarding an e-commerce site. The algorithm described below combines the labels from these Label Bureaus into an overall rating system for the consumer.
- **Filtering.** PICSRules are applied to this overall rating to determine if the e-commerce site meets the consumer's criteria.
- **Ranking.** Those companies that pass the filter can be presented in a rank order dependent on their derived ratings.

The architecture of the system is shown in Figure 3. The third-party labels are stored in Label Bureaus distributed over the Internet. The Label Engine maintains a data base of these Label Bureaus. In processing a client request, the following steps take place:

- The Application Server receives a request from a client machine for information about an e-commerce Web site,
- The Application Server issues a request to the Label Engine for the list of Label Bureaus that hold labels rating the e-commerce of interest,
- The Label Engine issues a request to all of the Label Bureaus in its data base asking if they hold labels for this e-commerce site,
- The Label Bureaus return a simple yes/no to the Label Engine,
- The Label Engine then returns the URLs of those Label Bureaus holding such labels to the Application Server,
- The Application Server then makes requests for those labels directly to the Label Bureaus and requests the associated rating systems from the Rating Services
- Upon receiving the labels and the rating systems, the Application Server combines these labels and their associated rating systems to produce an overall rating for the e-commerce site and, possibly, to filter out the site.

This is a two-step process in that the Application Server asks the Label Engine to find out which Label Bureaus hold the appropriate labels, and then the Application Server goes back to the Label Bureaus to retrieve the labels. Although the Label Engine could return the labels directly to the Application Server, this design permits the Application Server to bypass the Label Engine to access a known Label Bureau directly. It also takes a load off the Label Engine in anticipation of such an engine becoming heavily used.

Once the labels have been retrieved, the associated rating systems are retrieved from the Rating Services. Each rating system specifies the dimensions used for labeling, scale of allowable values on each dimension, and a descriptor of the criteria used in assigning values [11]. The Application Server combines these labels and their associated rating systems to produce an overall rating.

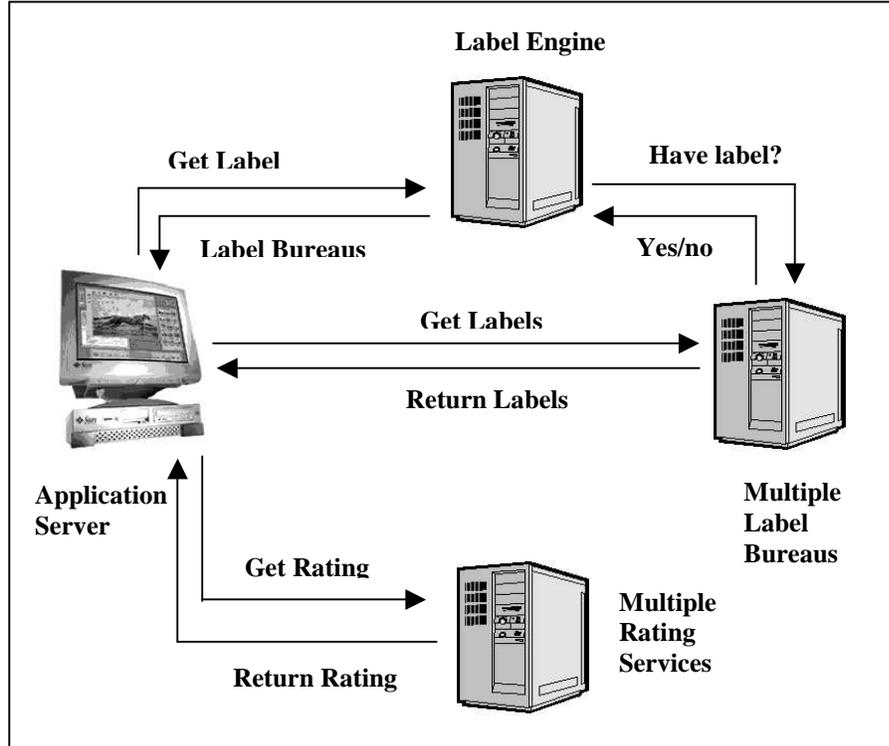


Figure 3. Architecture

4.1 PICSRules in Context of this Architecture:

PICSRules, although commonly used for safe Internet browsing and content filtering, can be used to describe e-commerce sites and services. In this architecture, PICSRules are used for the following purposes:

- List Rating Services and Label Bureaus
- Set user preferences for computing the ratings. Select whether to rate, filter or rank the site(s)
- Select the rating algorithm, if the server allows multiple rating algorithms for computing the combined rating.
- Select a specific dimension, if the user decides to get the rating on that particular category.
- Optionally set the weights of the dimensions, if the user is not satisfied with the default weights.

It would be easier to explain and understand the above mentioned features of the PICSRules with an example from our prototype system. Consider a sample PICSRule:

(PicsRule-1.1

```
(
  ServiceInfo (
    'http://borg.cs.dal.ca:2084/1.rat'
    shortname "Rat"
    bureauURL 'http://borg.cs.dal.ca:2084'
  )
  reqextension (
    'http://borg.cs.dal.ca:2086/ratOption.html'
    shortname 'ratOption'
  )
  ratOption (
    option 'RATING'
    specifyBureau 'YES'
    specifyEngine 'YES'
    ratAlgorithm 'DEFAULT'
  )
  reqextension (
    'http://borg.cs.dal.ca:2085/engineExtension.html'
    shortname "engineEntension"
  )
  engineExtension (
    address 'http://borg.cs.dal.ca'
    port '2081'
    UnAvaliable 'FAIL'
    MultipleEngine 'NO'
  )
)
)
```

In this example, only labels retrieved from the Label Bureau site 'http://borg.cs.dal.ca:2084' using the rating scheme 'http://borg.cs.dal.ca:2084/1.rat' are used to assess the e-commerce site. This is specified in the *ServiceInfo* portion of the rule.

ServiceInfo is used to specify a user selected Label Bureau and Rating Scheme for use in the evaluation. In such a case, even if the Label Engine does not include this Label Bureau in its list, the rating process will access their labels. So the labels from this Label Bureau will be procured along with the Label Engine specified Label Bureaus.

This rule also has two required extensions (*reqextension*). These extensions permit the consumer to specify certain options (*ratOption*) and to specify the Label Engine to the Application Server (*engineExtension*). These must be included in the rating process. If there is any error while retrieving these extensions, then the rating process fails, i.e., returns an error to the user.

ratOption is a sample extension that allows the user to choose one of the three options: rating, filtering or ranking. This is declared in the option dimension of the extension. Here the *specifyBureau* is the user specified Label Bureau to be included in the evaluation. The user has also specified the Label Engine to be used in this rule, hence the *specifyEngine* dimension. However the user wants the Application Server to use its default rating algorithm.

engineExtension is an extension created to specify the Label Engine details such as address and port number. If the Label Engine is unavailable, the rating process should terminate and return an error. This is specified by the *UnAvailable* option. The user desires to use only one engine to obtain the rating; hence the specified *multipleEngine* option as 'NO'.

5. The Default Rating Algorithm

The Application Server is responsible for applying the rating algorithm to the collected labels and rating systems. The rating algorithm provides the logic for processing labels in order to compute a rating value for the company associated with the given URL. The architecture is not based on a specific rating algorithm, but is designed in such a way that any rating algorithm can be incorporated.

The default rating algorithm permits the combination of different rating systems, with the user specifying the weight to be assigned to the values in any given dimension of a label. For example, in Figure 4, the default rating algorithm is applied to determine an overall evaluation of a particular e-commerce site. In this example, there are labels from two rating systems, A and B, to be combined. The dimension names and values are given in the boxes, with the security dimension of rating system A and the dispute resolution dimension of rating system B having sub-dimensions. All dimension weights are on the 0-10 scale with 10 being good. The user has assigned proportionate weights to each dimension, as shown in parentheses.

Weighted averages are calculated from the leaves to the root of the tree to determine the final result. In Figure 4, the value of each leaf node is assigned as shown and the value of each interior node is calculated as follows:

- value of the security dimension of scheme A is calculated as $(5*.6 + 7*.4) = 5.8$
- value of rating system A is calculated as $(5.8*.8 + 6*.2) = 5.84$
- value at dispute resolution dimension of scheme B is calculated as $(6*.7 + 3*.3) = 5.1$
- value of rating system B is calculated as $(6*.6 + 5.1*.4) = 5.64$

- rating systems A and B contribute equal amounts to the final result value of for this e-commerce site and calculated as: $(5.84 \cdot .5 + 5.64 \cdot .5) = 5.74$

Further PICSRules would then be applied in order to determine whether or not to accept or to filter out this e-commerce site.

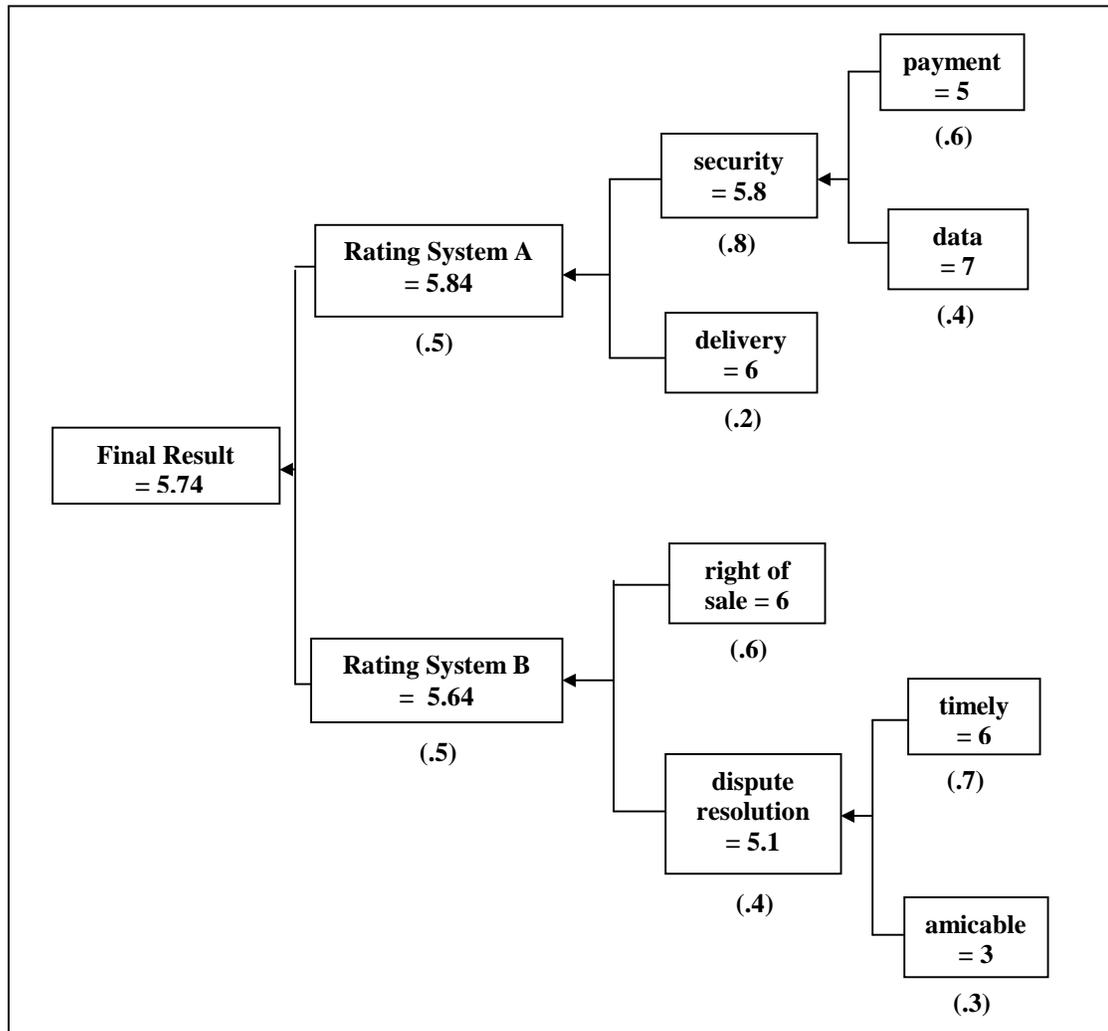


Figure 4. Default Rating Algorithm

In order to allow the user to specify the rating algorithm, we make use of the extension clause of PICSRules. In the above PICSRule example, this extension is termed as ‘ratOption’ with the dimension ratAlgorithm specifying the name of the rating algorithm. In the above example, a default ranking algorithm was applied.

However, if the user wishes to compute the rating on the basis of a particular dimension, the user can mention it in the PICSRules itself, using the extension clause termed ‘dimensionExtension’ to specify the dimension. The nomenclature of extension is left for individual implementations.

If the dimension is mentioned, then only labels for those Rating Systems having that dimension are requested and the rating is calculated for that particular dimension only. For example in our rating algorithm, we would bubble the values up the dimension tree only till that dimension is reached, then use this rating as final for that dimension tree. If multiple labels are available for that dimension, the rating algorithm treats these dimension sub-trees as rating trees and combines their ratings.

These user-specified weights can be specified in an extension clause, which is interpreted by the Application Server. It is important to note that these weights are used in our rating algorithm, which is not a mandatory part of the architecture. Software has also been developed for creating the PICSRules and labels and for the retrieval and viewing of the e-commerce sites, including the labels and rules.

6. Summary

In summary, we have developed a prototype architecture that permits distributed Label Bureaus and multiple rating systems to be combined in such a fashion that users can rate, filter and rank e-commerce sites on the Web. PICS and PICSRules are used as the platform for the implementation. The rating can be done either by specific category or by a default algorithm that considers all dimensions. The architecture is sufficiently flexible to permit various rating algorithms to be applied.

From the perspective of the customer, the ‘legitimacy’ conditions that are required for the development of trust for B2C e-commerce include:

- the sellers are who they claim to be
- the seller has right of sale over the item in question
- the transaction and payment mechanisms are available, legal and secure
- information about the buyer is not redistributed to other organizations or used for other purposes than for which it was intended
- the item sold corresponds to its description and is suitable for its intended purpose
- the purchased item can and will be delivered to the buyer

The prototype system, as briefly described in this paper, is an e-commerce support service that can help create these legitimacy conditions. PICS labels supports multi-dimensional rating schemes, thus permitting e-commerce sites to be rated along multiple dimensions and allowing consumers to have access to these ratings. These dimensions may include all of the above legitimacy conditions and may be arranged in a hierarchical manner. The layer cake model permits different third-party templates to be applied, allowing for differences in countries and cultures. Most importantly, the consumer can generate their own PICSRules specifying how they wish labels to be processed. This

includes the ability to combine labels generated by different rating services using different rating schema.

It is proposed that rating systems be developed along these dimension lines and that trusted third-party rating be done and the resulting labels stored in Label Bureaus. If such labeling is done, then a distributed network of Label Bureaus with multiple rating schemes and some architecture, such as we have prototyped, can be used to access and use this information effectively.

7. References

- [1] J.M. Balkin, B.S. Noveck and K. Roosevelt. (1999). Filtering the internet: A best practices model. Information and Society Project. Available: http://www.law.yale.edu/infosociety/filtering_report.html [2000, December]
- [2] Bertelsmann Foundation. Self Regulation of Internet Content. Available: http://www.stiftung.Bertelsmann.de/internetcontent/english/frameset_home.htm [2000, December]
- [3] Better Business Bureau (Online). Available: <http://www.BBBonline.org> [2000, December]
- [4] eBay. Available: <http://www.eBay.com> [2000, December]
- [5] Christopher Evans, Clive D.W. Feather, Alex Hopmann, Martin Presler-Marshall and Paul Resnick. (1997). PICS Rules 1.1, W3C Recommendation 29 Dec 1997 Available: <http://www.w3.org/TR/REC-PICSRules> [2000, December]
- [6] G. Eysenback. (1999). Collaboration for Critical Appraisal of Medical Information on the Internet. Available: <http://www.dermis.net/medpics/> [2000, December]
- [7] Sam Kinney. (2000). An Overview of B2B and Purchasing Technology. Response to Call for Submissions, Federal Trade Commission, Public Workshop: Competition Policy In the World of B2B Electronic Marketplaces. Available: http://mba.tuck.Dartmouth.edu/Anderson/Kinney_white_paper.htm. [2000, December]
- [9] Tim Krauskopf, Jim Miller, Paul Resnick and Win Treese. (1996). PICS Label Distribution Label Syntax and Communication Protocols, Version 1.1, W3C Recommendation 31-October-96. Available: <http://www.w3.org/TR/REC-PICS-labels> [2000, December]
- [10] Jayson Matthews. (2000). Study Quantifies B2C Problems. *InternetNews*. August 1, 2000. Available: http://www.internetnews.com/bus-news/articles/0,,3_427751,00.html. [2000, December]

- [11] Jim Miller, Paul Resnick and David Singer. (1996). Rating Services and Rating Systems (and Their Machine Readable Descriptions), Version 1.1, W3C Recommendations 31 Oct 1996. Available: <http://www.w3c.org/TR/REC-PICS-services-961031> [2000, December]
- [12] RSACi Recreational Software Advisory Council. 2000. Available: <http://www.icra.org>. [2000, January]
- [13] SafeSurf. Available: <http://www.safesurf.com>. [2000, January]
- [14] M. Shepherd and C. Watters. (2000). Content Filtering Technologies and Internet Service Providers: Enabling User Choice, Industry Canada. Canada. March, 2000. Available: <http://strategis.ic.gc.ca/internet>. [2000, December]
- [15] Koh Ai Tee. (2000). E-Commerce in an Era of Creative Destruction. Available: <http://www.alumni.nus.edu.sg/Almnus/jul2000/ecom.html>. [2000, December]
- [16] TRUSTe. Available: <http://www.Truste.com> [2000, December]