



White Paper

RULES BASED CONTENT DELIVERY AND THE TAILORED USER EXPERIENCE

Table of Contents

Introduction 3
Classifying The Content 3
Delivery Models 4
Content Provisioning 4
Searchable Content 4
Customized Dynamic Navigation 5
Rules and Representations 5
Representing Rules as Content 6
The SR2 Example 6
Conclusion 7





Introduction

What is it that makes information so easy to accumulate, yet so difficult to disseminate? The Internet, for all of its accessibility, convenience, and speed, is not very structured. It has no keeper (OK Google tries), no librarian to organize, sort, and classify information, no guardian to separate the fiction from the factual. It contains vast quantities of data, in the form of statistics, reports, essays, and books, but it also contains an amazing amount of irrelevance and misinformation (any Google search will prove that one out) unfortunately this also applies to most corporate websites today. To effectively communicate via the Internet, information needs to be locatable by the people who want it, when they want it.

Each enterprise has to determine the best ways to make their information not only accurate and relevant, but easily delivered to an interested audience.

Classifying the Content

In order to achieve the goal of intelligent, well-delivered content, it is imperative that the content itself be organized in ways that make sense to the common audience. Classification can be done in many ways. Selecting the best form of classification can be an involved process, requiring extensive analysis and multiple iterations. In fact, the process of organizing information and presenting it within the framework of the organization is significant enough to have developed into an area of specialized study, commonly referred to as Information Architecture (IA), and a full discussion of the concepts, methods and tools of this discipline can be found in our Refresh Software Information Architecture white paper.

The first and perhaps most prominent method of classification today is through a hierarchical taxonomy. This sort of classification scheme is modeled around multiple layers of groupings and sub-groupings of artifacts, where the categories of each successive layer provide a more precise definition of its members.

Another modern technique increasingly used to classify content is faceted or aspect oriented classification. With this technique, information is tagged with values describing several relevant attributes, and users can find data by using combinations of these attributes to narrow in on the content they need to locate, based on what is important to the user in their current search. Unlike the hierarchical approach, the content is classified in flat categories, but across multiple axes of metadata.

Often used in conjunction with the faceted approach, is the use of controlled vocabularies for increased precision. A controlled vocabulary is a set of precisely defined, and agreed upon terms about a subject area, which can be used as classification and search terms. Such lists are often

developed jointly by experts in a subject, and published for reference by others.

Classification of data, however important, is still only part of the picture. To actively and intelligently deliver relevant and timely content to interested users, the system must be given information about the users, as well as about the content available. In some environments, information is guaranteed to be known about users, because the system itself is private, so access is only granted after some minimum level of information on a user is gathered and entered into the system. In more public environments, voluntary registration, user activity tracking, and other means are used to understand the requirements and interests of users.



Delivery Models

Once both the data to be delivered, and the users looking for it, have been classified, there is still the question of how it is delivered. There are several strategies, each with advantages and disadvantages, which can be used to deliver the content to a user.

Content Provisioning

Content provisioning refers to placing content into place, ready to be accessed by one or more users. In the web space, this might mean a showcase area on the homepage for newly available and/or recently updated content, or an RSS feed, or similar syndication mechanism. In the mobile arena, it might mean notification of availability, along with the opportunity to retrieve content, at the user's convenience. For print media, such as a product catalog, it would mean starting the process of printing and mailing out the actual physical catalog.

Intelligent provisioning would not merely deliver the content to the various users, but it would deliver it in a preferred medium first, with highlights of relevant content items based on user profiling and content categorization. It would tailor not merely the selection of content, but the organization and prominence of each item, based on the user profile. Different mediums could use different thresholds for provisioning, resulting in wireless provisioning of only the highest items, web-based provisioning of several more, and email notification, or print distribution of the full set of content, organized by relevance and interest.

Searchable Content

Sometimes, the best way to deliver the right content to a user is to wait until the user asks for it. And, making content searchable means more than merely using a full text search engine. A well-classified and organized body of information can be made vastly more accessible to the common user, by exposing some of the classification mechanisms

employed within it. For example, if controlled vocabularies have been employed as part of the classification scheme, then using those terms to narrow a search will produce far better results for the user. Providing an interface which suggests alternate terms, based on matches within the vocabulary itself, can help a user locate what they need, without requiring them to be fluent in the industry jargon that may comprise the vocabulary.



Customized Dynamic Navigation

Another approach to delivering content is to correlate similar user interests, and cross-link directly to items heavily visited by those with similar profiles. This tactic is very similar to one often seen in websites devoted to product sales, reflecting items often purchased together, on the assumption that customers buying the same merchandise probably have some level of shared interest in related topics. Providing direct access to content with a probability of high relevance to the user, based on their profile and current content selection provides a very customized experience

Rules and Representations

The final piece missing from the puzzle; however, is perhaps the most interesting. None of the magic of a highly dynamic, personally customized system happens on its own. At the most elementary level, the delivery system can represent the rules directly into the software driving the delivery. The vast majority of content delivery systems, today, do exactly that. The business rules defining which users should see what data based on profiling, security concerns, customer status, and so on, are turned into code, and used to drive the content delivery. However, when using this approach, when those rules change the code has to change. That means developer time, a release schedule, and a loss of agility in the marketplace - agility that, in today's world of high speed business and even higher speed marketing, can mean the difference between owning the space, and scrambling for the scraps.

With a little ingenuity, a delivery system can be constructed that determines its rules from sets of configurations - mappings between user profile attributes and content facets; security constraints; ranking algorithms; decision criteria for distribution mediums - and then acts according to the configured plan. Such a system might contain support for rule chaining, vetoing, and prioritization, allowing for complex systems of rules to be built up governing the delivery of content through the system. Unfortunately, managing the configuration of such a system quickly becomes cumbersome, and may easily grow into a task too complicated and arcane for the real decision makers about content delivery to maintain without help.

The technologically mature approach resolves this issue, as well, by modeling the rules as content. When the mappings between user interests, and content subjects are maintained alongside the content itself, and processed in real time, the true power and agility of a mature Content Management Solution can finally be realized.

Representing Rules as Content

The key to this solution lies in the ability to treat the rules defining the delivery strategy as content artifacts themselves. Not all Content Management Systems can do this, and even fewer can do it well. Many content systems, especially those which are specifically geared for the web, are 'page-oriented', defining content by the page it appears on. Such systems are clearly ill-suited for representing data that is never displayed to an end-user. Others are 'content-oriented', allowing sections of pages to be separated, managed independently, and reassembled for presentation, but often failing to support a clean separation between the content, and the presentation of the content, resulting in content units which are difficult to present across different mediums, and which often cannot gracefully handle non-viewable data. More flexible Content Management Systems like SR2 Component Content Management are 'object oriented' or 'asset oriented' Content Management Systems, are designed to capture structured units of content, using discrete fields bundled together as a 'content asset' or 'content object'. These units of content are capable of capturing not only the viewable content, but the invisible metadata assignments, whether as facets, or as taxonomy classifications, and more. If the system is built right, they can capture these metadata blocks, and build relationships between them, even devoid of viewable content.

By capturing the classifications assigned to a unit of content as part of the content definition, the system is able to maintain the classifications, and even allow editors to reclassify content when needed. Further, when the classifications themselves are maintained as content, new categories can be introduced seamlessly, changes in structure (such as a new layer added to a taxonomy) can take place at the content management level, and relationships between classifications can be maintained.

The SR2 Component CMS Example

Implementing a system, using a robust and mature asset oriented Content Management System (CMS) such as Refresh Software SR2 CMS is straightforward. The Asset Types (or, in object-oriented terms object types or classes) used for each group of content are built to provide fields for the relevant metadata used in classification. In this example case, a product asset type has fields for region, and channel. As these fields have limited, specific selection, a multi-select control would be the optimal choice for both. SR2 multi-select fields can draw the input options from database tables, either static lookup tables, or a table of assets built by



the system. In this case, we use a Region asset type to represent a region rule, and drive the multi-select field from that. Similarly, the channel field could be driven by a channel asset.

The region Asset Type, itself, besides naming the region, would contain its own multi-select of states, which, because it is unlikely to change, we populate from a standardized state table rather than an asset list. In a complete system, it likely also contains additional information about the region, such as the regional sales manager, support contact information, and regional office location.

The users are maintained in a separate system, and checked in real-time by the delivery system. For the purposes of this example, it only needs to maintain the state in which the dealer does business, and the nature of the business relationship, information which can reasonably be assumed to be in any typical Customer Relationship Management (CRM) database.

All that remains is some straightforward code, which contains the logic to follow these rules and mappings, to deliver the content where it belongs. Since the relationships between the various artifacts are maintained as part of the content itself, the content retrieval APIs provided by SR2 provide simple means to perform filtered lookups at the point of delivery, to do just that. Because SR2 also provides an open schema, with full visibility into the database, this could even be written as one or more stored procedures, for delivery into other database driven applications.

Conclusion

Intelligent content delivery does not happen without forethought and planning. Content classification, user definition, delivery rules, and content relationships can be nebulous to define, and difficult to agree on within an enterprise environment. However, the benefits of an agile implementation for content delivery clearly show it is effort well spent. An underlying content management tool with proper support for asset oriented content management, such as Refresh Software's SR2 Component CMS, can make the implementation easier, faster, and more robust. But the benefits don't stop there. By putting the control of the content delivery directly into the hands of the decision makers, an enterprise can maintain the market agility needed to thrive in the modern marketplace.

